

Game theory

Contents

1	Introduction	2
2	Combinatorial games	7
2.1	Some definitions	7
2.2	The game of nim, and Bouton's solution	10
2.3	The sum of combinatorial games	14
2.4	Staircase nim and other examples	18
2.5	The game of Green Hackenbush	20
2.6	Wythoff's nim	21
3	Two-person zero-sum games	23
3.1	Some examples	23
3.2	The technique of domination	25
3.3	The use of symmetry	27
3.4	von Neumann's minimax theorem	28
3.5	Resistor networks and troll games	31
3.6	Hide-and-seek games	33
3.7	General hide-and-seek games	34
3.8	The bomber and submarine game	37
3.9	A further example	38
4	General sum games	39
4.1	Some examples	39
4.2	Nash equilibrium	40
4.3	General sum games with $k \geq 2$ players	44
4.4	The proof of Nash's theorem	45
4.4.1	Some more fixed point theorems	47
4.4.2	Sperner's lemma	49
4.4.3	Proof of Brouwer's fixed point theorem	51
4.5	Some further examples	51
4.6	Potential games	52

VISIT...

LANZAROTE
Caliente.COM

5	Coalitions and Shapley value	55
5.1	The Shapley value and the glove market	55
5.2	Probabilistic interpretation of Shapley value	57
5.3	Two more examples	59
6	Mechanism design	61

1 Introduction

In this course on game theory, we will be studying a range of mathematical models of conflict and cooperation between two or more agents. The course will attempt an overview of a broad range of models that are studied in game theory, and that have found application in, for example, economics and evolutionary biology. In this Introduction, we outline the content of this course, often giving examples.

One class of games that we begin studying are combinatorial games. An example of a combinatorial game is that of **hex**, which is played on an hexagonal grid shaped as a rhombus: think of a large rhombus-shaped region that is tiled by a grid of small hexagons. Two players, R and G, alternately color in hexagons of their choice either red or green, the red player aiming to produce a red crossing from left to right in the rhombus and the green player aiming to form a green one from top to bottom. As we will see, the first player has a winning strategy; however, finding this strategy remains an unsolved problem, except when the size of the board is small (9×9 , at most). An interesting variant of the game is that in which, instead of taking turns to play, a coin is tossed at each turn, so that each player plays the next turn with probability one half. In this variant, the optimal strategy for either player is known.

A second example which is simpler to analyse is the game of **nim**. There are two players, and several piles of sticks at the start of the game. The players take turns, and at each turn, must remove at least one stick from one pile. The player can remove any number of sticks that he pleases, but these must be drawn from a single pile. The aim of the game is to force the opponent to take the last stick remaining in the game. We will find the solution to nim: it is not one of the harder examples.

Another class of games are **congestion games**. Imagine two drivers, *I* and *II*, who aim to travel from cities *B* to *D*, and from *A* to *C*, respectively:

A	(1,2)	D
(3,5)		(2,4)
B	(3,4)	C

The costs incurred to the drivers depend on whether they travel the roads alone or together with the other driver (not necessarily at the very same time). The vectors (a, b) attached to each road mean that the cost paid by either driver for the use of the road is a if he travels the road alone, and b if he shares its use with the other driver. For example, if I and II use the road AB — which means that I chooses the route via A and II chooses that via B — then each pays 5 units for doing so, whereas if only one of them uses that road, the cost is 3 units to that driver. We write a cost matrix to describe the game:

	II	B	D
I			
A	(6,8)	(5,4)	
C	(6,7)	(7,5)	

The vector notation $(\cdot, \cdot)$ denotes the costs to players I and II of their joint choice.

A fourth example is that of **penalty kicks**, in which there are two participants, the penalty-taker and the goalkeeper. The notion of left and right will be from the perspective of the goalkeeper, not the penalty-taker. The penalty-taker chooses to hit the ball either to the left or the right, and the goalkeeper dives in one of these directions. We display the probabilities that the penalty is scored in the following table:

	GK	L	R
PT			
L	0.8	1	
R	1	0.5	

That is, if the goalie makes the wrong choice, he has no chance of saving the goal. The penalty-taker has a strong ‘left’ foot, and has a better chance if he plays left. The goalkeeper aims to minimize the probability of the penalty being scored, and the penalty-taker aims to maximize it. We could write a payoff matrix for the game, as we did in the previous example, but, since it is zero-sum, with the interests of the players being diametrically opposed, doing so is redundant. We will determine the optimal strategy for the players for a class of games that include this one. This strategy will often turn out to be a randomized choice among the available options.

Such **two person zero-sum games** have been applied in a lot of contexts: in sports, like this example, in military contexts, in economic applications, and in evolutionary biology. These games have a quite complete theory, so that it has been tempting to try to apply them. However, real life is often more complicated, with the possibility of cooperation between players to realize a mutual advantage. The theory of games that model such an effect is much less complete.

The mathematics associated to zero-sum games is that of convex geometry. A convex set is one where, for any two points in the set, the straight line segment connecting the two points is itself contained in the set.

The relevant geometric fact for this aspect of game theory is that, given any closed convex set in the plane and a point lying outside of it, we can find a line that separates the set from the point. There is an analogous statement in higher dimensions. von Neumann exploited this fact to solve zero sum games using a **minimax** variational principle. We will prove this result.

In **general-sum games**, we do not have a pair of optimal strategies any more, but a concept related to the von Neumann minimax is that of **Nash equilibrium**: is there a ‘rational’ choice for the two players, and if so, what could it be? The meaning of ‘rational’ here and in many contexts is a valid subject for discussion. There are anyway often many Nash equilibria and further criteria are required to pick out relevant ones.

A development of the last twenty years that we will discuss is the application of game theory to evolutionary biology. In economic applications, it is often assumed that the agents are acting ‘rationally’, and a neat theorem should not distract us from remembering that this can be a hazardous assumption. In some biological applications, we can however see Nash equilibria arising as stable points of evolutionary systems composed of agents who are ‘just doing their own thing’, without needing to be ‘rational’.

Let us introduce another geometrical tool. Although from its statement, it is not evident what the connection of this result to game theory might be, we will see that the theorem is of central importance in proving the existence of Nash equilibria.

Theorem 1 (Brouwer’s fixed point theorem) : *If $K \subseteq \mathbb{R}^d$ is closed, bounded and convex, and $T : K \rightarrow K$ is continuous, then T has a fixed point. That is, there exists $x \in K$ for which $T(x) = x$.*

The assumption of convexity can be weakened, but not discarded entirely. To see this, consider the example of the annulus $C = \{x \in \mathbb{R}^2 : 1 \leq |x| \leq 2\}$, and the mapping $T : C \rightarrow C$ that sends each point to its rotation by 90 degrees anticlockwise about the origin. Then T is *isometric*, that is, $|T(x) - T(y)| = |x - y|$ for each pair of points $x, y \in C$. Certainly then, T is continuous, but it has no fixed point.

Another interesting topic is that of signalling. If one player has some information that another does not, that may be to his advantage. But if he plays differently, might he give away what he knows, thereby removing this advantage?

A quick mention of other topics, related to **mechanism design**. Firstly, voting. Arrow's impossibility theorem states roughly that if there is an election with more than two candidates, then no matter which system one chooses to use for voting, there is trouble ahead: at least one desirable property that we might wish for the election will be violated. A recent topic is that of eliciting truth. In an ordinary auction, there is a temptation to underbid. For example, if a bidder values an item at 100 dollars, then he has no motive to bid any more or even that much, because by exchanging 100 dollars for the object at stake, he has gained an item only of the same value to him as his money. The second-price auction is an attempt to overcome this flaw: in this scheme, the lot goes to the highest bidder, but at the price offered by the second-highest bidder. This problem and its solutions are relevant to bandwidth auctions made by governments to cellular phone companies.

Example: Pie cutting. As another example, consider the problem of a pie, different parts of whose interior are composed of different ingredients. The game has two or more players, who each have their own preferences regarding which parts of the pie they would most like to have. If there are just two players, there is a well-known method for dividing the pie: one splits it into two halves, and the other chooses which he would like. Each obtains at least one-half of the pie, as measured according to each own preferences. But what if there are three or more players? We will study this question, and a variant where we also require that the pie be cut in such a way that each player judges that he gets at least as much as anyone else, according to his own criterion.

Example: Secret sharing. Suppose that we plan to give a secret to two people. We do not trust either of them entirely, but want the secret to be known to each of them provided that they co-operate. If we look for a physical solution to this problem, we might just put the secret in a room, put two locks on the door, and give each of the players the key to one of the locks. In a computing context, we might take a password and split it in two, giving each half to one of the players. However, this would force the length of the password to be high, if one or other half is not to be guessed by repeated tries. A more ambitious goal is to split the secret in two in such a way that neither person has any useful information on his own. And here is how to do it: suppose that the secret s is an integer that lies between 0 and some large value M , for example, $M = 10^6$. We who hold the secret at the start produce a random integer x , whose distribution is uniform on the interval $\{0, \dots, M - 1\}$ (*uniform* means that each of the M possible

outcomes is equally likely, having probability $1/M$). We tell the number x to the first person, and the number $y = (s - x) \bmod M$ to the second person ($\bmod M$ means adding the right multiple of M so that the value lies on the interval $\{0, \dots, M - 1\}$). The first person has no useful information. What about the second? Note that

$$\mathbb{P}(y = j) = \mathbb{P}((s - x) \bmod M = j) = 1/M,$$

where the last equality holds because $(s - x) \bmod M$ equals y if and only if the uniform random variable x happens to hit one particular value on $\{0, \dots, M - 1\}$. So the second person himself only has a uniform random variable, and, thus, no useful information. Together, however, the players can add the values they have been given, reduce the answer $\bmod M$, and get the secret s back. A variant of this scheme can work with any number of players. We can have ten of them, and arrange a way that any nine of them have no useful information even if they pool their resources, but the ten together can unlock the secret.

Example: Cooperative games. These games deal with the formation of coalitions, and their mathematical solution involves the notion of **Shapley value**. As an example, suppose that three people, *I*, *II* and *III*, sit in a store, the first two bearing a left-handed glove, while the third has a right-handed one. A wealthy tourist, ignorant of the bitter local climatic conditions, enters the store in dire need of a pair of gloves. She refuses to deal with the glove-bearers individually, so that it becomes their job to form coalitions to make a sale of a left and right-handed glove to her. The third player has an advantage, because his commodity is in scarcer supply. This means that he should be able to obtain a higher fraction of the payment that the tourist makes than either of the other players. However, if he holds out for too high a fraction of the earnings, the other players may agree between them to refuse to deal with him at all, blocking any sale, and thereby risking his earnings. We will prove results in terms of the concept of the Shapley value that provide a solution to this type of problem.

2 Combinatorial games

2.1 Some definitions

Example. We begin with n chips in one pile. Players I and II make their moves alternately, with player I going first. Each player takes between one and four chips on his turn. The player who removes the last chip wins the game. We write

$$\mathbf{N} = \{n \in \mathbb{N} : \text{player } I \text{ wins if there are } n \text{ chips at the start}\},$$

where we are assuming that each player plays optimally. Furthermore,

$$\mathbf{P} = \{n \in \mathbb{N} : \text{player } II \text{ wins if there are } n \text{ chips at the start}\}.$$

Clearly, $\{1, 2, 3, 4\} \subseteq \mathbf{N}$, because player I can win with his first move. Then $5 \in \mathbf{P}$, because the number of chips after the first move must lie in the set $\{1, 2, 3, 4\}$. That $\{6, 7, 8, 9\} \in \mathbf{N}$ follows from the fact that player I can force his opponent into a losing position by ensuring that there are five chips at the end of his first turn. Continuing this line of argument, we find that $\mathbf{P} = \{n \in \mathbb{N} : n \text{ is divisible by five}\}$.

Definition 1 A **combinatorial game** has two players, and a set, which is usually finite, of possible positions. There are rules for each of the players that specify the available legal moves for the player whose turn it is. If the moves are the same for each of the players, the game is called **impartial**. Otherwise, it is called **partisan**. The players alternate moves. Under **normal** play, the player who cannot move loses. Under **misère** play, the player who makes the final move loses.

Definition 2 Generalising the earlier example, we write $\mathbf{N}$ for the collection of positions from which the next player to move will win, and $\mathbf{P}$ for the positions for which the other player will win, provided that each of the players adopts an optimal strategy.

Writing this more formally, assuming that the game is conducted under normal play, we define

$$\begin{aligned} \mathbf{P}_0 &= \{0\} \\ \mathbf{N}_{i+1} &= \{\text{positions } x \text{ for which there is a move leading to } \mathbf{P}_i\} \\ \mathbf{P}_i &= \{\text{positions } y \text{ such that each move leads to } \mathbf{N}_i\} \end{aligned}$$

for each $i \in \mathbb{N}$. We set

$$\mathbf{N} = \bigcup_{i \geq 0} \mathbf{N}_i, \quad \mathbf{P} = \bigcup_{i \geq 0} \mathbf{P}_i.$$


A strategy is just a function assigning a legal move to each possible position. Now, there is the natural question whether all positions of a game lie in $\mathbf{N} \cup \mathbf{P}$, i.e., if there is a winning strategy for either player.

Example: hex. Recall the description of hex from the Introduction, with R being player *I*, and G being player *II*. This is a partisan combinatorial game under normal play, with terminal positions being the colorings that have either type of crossing. (Formally, we could make the game “impartial” by letting both players use both colors, but then we have to declare two types of terminal positions, according to the color of the crossing.)


Note that, instead of a rhombus board with the four sides colored in the standard way, the game is possible to define on an arbitrary board, with a fixed subset of pre-colored hexagons — provided the board has the property that in any coloring of all its unfixed hexagons, there is exactly one type of crossing between the pre-colored red and green parts. Such pre-colored boards will be called admissible.

However, we have not even proved yet that the standard rhombus board is admissible. That there cannot be both types of crossing looks completely obvious, until you actually try to prove it carefully. This statement is the discrete analog of the Jordan curve theorem, saying that a continuous closed curve in the plane divides the plane into two connected components. This innocent claim has no simple proof, and, although the discrete version is easier, they are roughly equivalent. On the other hand, the claim that in any coloring of the board, there exists a monochromatic crossing, is the discrete analog of the 2-dimensional Brouwer fixed point theorem, which we have seen in the Introduction and will see proved in Section 4. The discrete versions of these theorems have the advantage that it might be possible to prove them by induction. Such an induction is done beautifully in the following proof, due to Craige Schensted.

Consider the **game of Y**: given a triangular board, tiled with hexagons, the two players take turns coloring hexagons as in hex, with the goal of establishing a chain that connects all three sides of the triangle.



Red has a winning Y here.




Reduction to hex.

Hex is a special case of Y: playing Y, started from the position shown on the right hand side picture, is equivalent to playing hex in the empty region of the board. Thus, if Y always has a winner, then this is also true for hex.

Theorem 2 *In any coloring of the triangular board, there is exactly one type of Y.*

Proof. We can reduce a colored board with sides of size n to a color board of size $n - 1$, as follows. Each little group of three adjacent hexagonal cells, forming a little triangle that is oriented the same way as the whole board, is replaced by a single cell. The color of the cell will be the majority of the colors of the three cells in the little triangle. This process can be continued to get a colored board of size $n - 2$, and so on, all the way down to a single cell. We claim that the color of this last cell is the color of the winner of Y on the original board.



Reducing a red Y to smaller and smaller ones.

Indeed, notice that any chain of connected red hexagons on a board of size n reduces to a connected red chain on the board of size $n - 1$. Moreover, if the chain touched a side of the original board, it also touches the side of the smaller one. The converse statement is just slightly harder to see: if there is a red chain touching a side of the smaller board, then there was a corresponding a red chain, touching the same side of the larger board. Since the single colored cell of the board of size 1 forms a winner Y on that board, there was a Y of the same color on the original board. $\square$

Going back to hex, it is easy to see by induction on the number of unfilled hexagons, that on any admissible board, one of the players has a winning strategy. One just has to observe that coloring red any one of the unfilled hexagons of an admissible board leads to a smaller admissible board, for which we can already use the induction hypothesis. There are two possibilities: (1) R can choose that first hexagon in such a way that on the resulting smaller board R has a winning strategy as being player II. Then R has a winning strategy on the original board. (2) There is no such hexagon, in which case G has a winning strategy on the original board.

Theorem 3 *On a standard symmetric hex board of arbitrary size, player I has a winning strategy.*

Proof. The idea of the proof is *strategy-stealing*. We know that one of the players has a winning strategy; suppose that player *II* is the one. This means that whatever player *I*'s first move is, player *II* can win the game from the resulting situation. But player *I* can pretend that he is player *II*: he just has to imagine that the colors are inverted, and that, before his first move, player *II* already had a move. Whatever move he imagines, he can win the game by the winning strategy stolen from player *II*; moreover, his actual situation is even better. Hence, in fact, player *I* has a winning strategy, a contradiction. $\square$

Now, we generalize some of the ideas appearing in the example of hex.

Definition 3 *A game is said to be **progressively bounded** if, for any starting position x , the game must finish within some finite number $B(x)$ of moves, no matter which moves the two players make.*

Example: Lasker's game. A position is finite collection of piles of chips. A player may remove chips from a given pile, or he may not remove chips, but instead break one pile into two, in any way that he pleases. To see that this game is progressively bounded, note that, if we define

$$B(x_1, \dots, x_k) = \sum_{i=1}^k (2x_i - 1),$$

then the sum equals the total number of chips and gaps between chips in a position $(x_1, \dots, x_k)$. It drops if the player removes a chip, but also if he breaks a pile, because, in that case, the number of gaps between chips drops by one. Hence, $B(x_1, \dots, x_k)$ is an upper bound on the number of steps that the game will take to finish from the starting position $(x_1, \dots, x_k)$.

Consider now a progressively bounded game, which, for simplicity, is assumed to be under normal play. We prove by induction on $B(x)$ that all positions lie in $\mathbf{N} \cup \mathbf{P}$. If $B(x) = 0$, this is true, because $\mathbf{P}_0 \subseteq \mathbf{P}$. Assume the inductive hypothesis for those positions x for which $B(x) \leq n$, and consider any position z satisfying $B(z) = n + 1$. There are two cases to handle: the first is that each move from z leads to a position in $\mathbf{N}$ (that is, to a member of one of the previously constructed sets $\mathbf{N}_i$). Then z lies in one of the sets $\mathbf{P}_i$ and thus in $\mathbf{P}$. In the second case, there is a move from z to some P -position. This implies that $z \in \mathbf{N}$. Thus, all positions lie in $\mathbf{N} \cup \mathbf{P}$.

2.2 The game of nim, and Bouton's solution

In the game of **nim**, there are several piles, each containing finitely many chips. A legal move is to remove any positive number of chips from a single pile. The aim of nim (under normal play) is to take the last stick remaining in the game. We will write the state of play in the game in the form

$(n_1, n_2, \dots, n_k)$, meaning that there are k piles of chips still in the game, and that the first has n_1 chips in it, the second n_2 , and so on.

Note that $(1, 1) \in \mathbf{P}$, because the game must end after the second turn from this beginning. We see that $(1, 2) \in \mathbf{N}$, because the first player can bring $(1, 2)$ to $(1, 1) \in \mathbf{P}$. Similarly, $(n, n) \in \mathbf{P}$ for $n \in \mathbb{N}$ and $(n, m) \in \mathbf{N}$ if $n, m \in \mathbb{N}$ are not equal. We see that $(1, 2, 3) \in \mathbf{P}$, because, whichever move the first player makes, the second can force there to be two piles of equal size. It follows immediately that $(1, 2, 3, 4) \in \mathbf{N}$. By dividing $(1, 2, 3, 4, 5)$ into two subgames, $(1, 2, 3) \in \mathbf{P}$ and $(4, 5) \in \mathbf{N}$, we get from the following lemma that it is in $\mathbf{N}$.

Lemma 1 *Take two nim positions, $A = (a_1, \dots, a_k)$ and $B = (b_1, \dots, b_\ell)$. Denote the position $(a_1, \dots, a_k, b_1, \dots, b_\ell)$ by (A, B) . If $A \in \mathbf{P}$ and $B \in \mathbf{N}$, then $(A, B) \in \mathbf{N}$. If $A, B \in \mathbf{P}$, then $(A, B) \in \mathbf{P}$. However, if $A, B \in \mathbf{N}$, then (A, B) can be either in $\mathbf{P}$ or in $\mathbf{N}$.*

Proof. If $A \in \mathbf{P}$ and $B \in \mathbf{N}$, then Player I can reduce B to a position $B' \in \mathbf{P}$, for which (A, B') is either terminal, and Player I won, or from which Player II can move only into pair of a $\mathbf{P}$ and an $\mathbf{N}$ -position. From that, Player I can again move into a pair of two $\mathbf{P}$ -positions, and so on. Therefore, Player I has a winning strategy.

If $A, B \in \mathbf{P}$, then any first move takes (A, B) to a pair of a $\mathbf{P}$ and an $\mathbf{N}$ -position, which is in $\mathbf{N}$, as we just saw. Hence Player II has a winning strategy for (A, B) .

We know already that the positions $(1, 2, 3, 4)$, $(1, 2, 3, 4, 5)$, $(5, 6)$ and $(6, 7)$ are all in $\mathbf{N}$. However, as the next exercise shows, $(1, 2, 3, 4, 5, 6) \in \mathbf{N}$ and $(1, 2, 3, 4, 5, 6, 7) \in \mathbf{P}$. $\square$

Exercise. By dividing the games into subgames, show that $(1, 2, 3, 4, 5, 6) \in \mathbf{N}$, and $(1, 2, 3, 4, 5, 6, 7) \in \mathbf{P}$. A hint for the latter one: adding two 1-chip piles does not affect the outcome of any position.

This divide-and-sum method still loses to the following ingenious theorem, giving a simple and very useful characterization of $\mathbf{N}$ and $\mathbf{P}$ for nim:

Theorem 4 (Bouton's Theorem) *Given a starting position $(n_1, \dots, n_k)$, write each n_i in binary form, and sum the k numbers in each of the digital places mod 2. The position is in $\mathbf{P}$ if and only if all of the sums are zero.*

To illustrate the theorem, consider the starting position $(1, 2, 3)$:

number of chips (decimal)	number of chips (binary)
1	01
2	10
3	11

Summing the two columns of the binary expansions modulo two, we obtain 00. The theorem confirms that $(1, 2, 3) \in \mathbf{P}$.

Proof of Bouton's Theorem. We write $n \oplus m$ to be the **nim-sum** of $n, m \in \mathbb{N}$. This operation is the one described in the statement of the theorem; i.e., we write n and m in binary, and compute the value of the sum of the digits in each column modulo 2. The result is the binary expression for the nim-sum $n \oplus m$. Equivalently, the nim-sum of a collection of values $(m_1, m_2, \dots, m_k)$ is the sum of all the powers of 2 that occurred an odd number of times when each of the numbers m_i is written as a sum of powers of 2. Here is an example: $m_1 = 13, m_2 = 9, m_3 = 3$. In powers of 2:

$$\begin{aligned} m_1 &= 2^3 + 2^2 && + 2^0 \\ m_2 &= 2^3 && + 2^0 \\ m_3 &= && + 2^1 + 2^0. \end{aligned}$$

In this case, the powers of 2 that appear an odd number of times are $2^0 = 1$ and $2^1 = 2$. This means that the nim-sum is $m_1 \oplus m_2 \oplus m_3 = 1 + 2 = 3$. For the case where $(m_1, m_2, m_3) = (5, 6, 15)$, we write, in purely binary notation,

5	0	1	0	1
6	0	1	1	0
15	1	1	1	1
	1	1	0	0

making the nim-sum 12 in this case. We define $\hat{P}$ to be those positions with nim-sum zero, and $\hat{N}$ to be all other positions. We claim that

$$\hat{P} = \mathbf{P} \quad \text{and} \quad \hat{N} = \mathbf{N}.$$

To check this claim, we need to show two things. Firstly, that $0 \in \hat{P}$, and that, for all $x \in \hat{N}$, there exists a move from x leading to $\hat{P}$. Secondly, that for every $y \in \hat{P}$, all moves from y lead to $\hat{N}$.

Note firstly that $0 \in \hat{P}$ is clear. Secondly, suppose that

$$x = (m_1, m_2, \dots, m_k) \in \hat{N}.$$

Set $s = m_1 \oplus \dots \oplus m_k$. Writing each m_i in binary, note that there are an odd number of values of $i \in \{1, \dots, k\}$ for which the binary expression for m_i has a 1 in the position of the left-most one in the expression for s . Choose one such i . Note that $m_i \oplus s < m_i$, because $m_i \oplus s$ has no 1 in this left-most position, and so is less than any number whose binary expression does have a 1 there. So we can play the move that removes from the i -th pile $m_i - m_i \oplus s$ chips, so that m_i becomes $m_i \oplus s$. The nim-sum of the resulting position $(m_1, \dots, m_{i-1}, m_i \oplus s, m_{i+1}, \dots, m_k)$ is zero, so this new

position lies in $\hat{P}$. We have checked the first of the two conditions which we require.

To verify the second condition, we have to show that if $y = (y_1, \dots, y_k) \in \hat{P}$, then any move from y leads to a position $z \in \hat{N}$. We write the y_i in binary:

$$\begin{aligned} y_1 &= y_1^{(n)} y_1^{(n-1)} \dots y_1^{(0)} = \sum_{j=0}^m y_1^{(j)} 2^j \\ &\dots \\ y_k &= y_k^{(n)} y_k^{(n-1)} \dots y_k^{(0)} = \sum_{j=0}^m y_k^{(j)} 2^j. \end{aligned}$$

By assumption, $y \in \hat{P}$. This means that the nim-sum $y_1^{(j)} \oplus \dots \oplus y_k^{(j)} = 0$ for each j . In other words, $\sum_{l=1}^k y_l^{(j)}$ is even for each j . Suppose that we remove chips from pile l . We get a new position $z = (z_1, \dots, z_k)$ with $z_i = y_i$ for $i \in \{1, \dots, k\}$, $i \neq l$, and with $z_l < y_l$. (The case where $z_l = 0$ is permitted.) Consider the binary expressions for y_l and z_l :

$$\begin{aligned} y_l &= y_l^{(n)} y_l^{(n-1)} \dots y_l^{(0)} \\ z_l &= z_l^{(n)} z_l^{(n-1)} \dots z_l^{(0)}. \end{aligned}$$

We scan these two rows of zeros and ones until we locate the first instance of a disagreement between them. In the column where it occurs, the nim-sum of y_l and z_l is one. This means that the nim-sum of $z = (z_1, \dots, z_k)$ is also equal to one in this column. Thus, $z \in \hat{N}$. We have checked the second condition that we needed, and so, the proof of the theorem is complete. $\square$

Example: the game of **rims**. In this game, a starting position consists of a finite number of dots in the plane, and a finite number of continuous loops. Each loop must not intersect itself, nor any of the other loops. Each loop must pass through at least one of the dots. It may pass through any number of them. A legal move for either of the two players consists of drawing a new loop, so that the new picture would be a legal starting position. The players' aim is to draw the last legal loop.

We can see that the game is identical to a variant of nim. For any given position, think of the dots that have no loop going through them as being divided into different classes. Each class consists of the set of dots that can be reached by a continuous path from a particular dot, without crossing any loop. We may think of each class of dots as being a pile of chips, like in nim. What then are the legal moves, expressed in these terms? Drawing a legal loop means removing at least one chip from a given pile, and then splitting the remaining chips in the pile into two separate piles. We can in fact split in any way we like, or leave the remaining chips in a single pile.

This means that the game of rims has some extra legal moves to those of nim. However, it turns out that these extra make no difference, and so that the sets **N** or **P** coincide for the two games. We now prove this.

Thinking of a position in rims as a finite number of piles, we write P_{nim} and N_{nim} for the $\mathbf{P}$ and $\mathbf{N}$ positions for the game of nim (so that these sets were found in Bouton's Theorem). We want to show that

$$\mathbf{P} = P_{nim} \quad \text{and} \quad \mathbf{N} = N_{nim}, \quad (1)$$

where $\mathbf{P}$ and $\mathbf{N}$ refer to the game of rims.

What must we check? Firstly, that $0 \in \mathbf{P}$, which is immediate. Secondly, that from any position in N_{nim} , we may move to P_{nim} by a move in rims. This is fine, because each nim move is legal in rims. Thirdly, that for any $y \in P_{nim}$, any rims move takes us to a position in N_{nim} . If the move does not involve breaking a pile, then it is a nim move, so this case is fine. We need then to consider a move where y_l is broken into two parts u and v whose sum satisfies $u + v < y$. Note that the nim-sum $u \oplus v$ of u and v is at most the ordinary sum $u + v$: this is because the nim-sum involves omitting certain powers of 2 from the expression for $u + v$. Thus,

$$u \oplus v \leq u + v < y_l.$$

So the rims move in question amounted to replacing the pile of size y_l by one with a smaller number of chips, $u \oplus v$. Thus, the rims move has the same effect as a legal move in nim, so that, when it is applied to $y \in P_{nim}$, it produces a position in N_{nim} . This is what we had to check, so we have finished proving (1).

Example: Wythoff nim. In this game, we have two piles. Legal moves are those of nim, but with the exception that it is also allowed to remove equal numbers of chips from each of the piles in a single move. This stops the positions $\{(n, n) : n \in \mathbb{N}\}$ from being $\mathbf{P}$ -positions. We will see that this game has an interesting structure.

2.3 The sum of combinatorial games

Definition 4 The **sum** of two combinatorial games, G_1 and G_2 , is that game G where, for any move, a player may choose in which of the games G_1 and G_2 to play. The terminal positions in G are (t_1, t_2) , where t_i is a terminal in G_i for both $i \in \{1, 2\}$. We will write $G = G_1 + G_2$.

We say that two pairs (G_i, x_i) , $i \in \{1, 2\}$, of a game and a starting position are **equivalent** if (x_1, x_2) is a $\mathbf{P}$ -position of the game $G_1 + G_2$. We will see that this notion of “equivalent” games defines an equivalence relation.

Optional exercise: Find a direct proof of transitivity of the relation “being equivalent games”.

As an example, we see that the nim position $(1, 3, 6)$ is equivalent to the nim position (4) , because the nim-sum of the sum game $(1, 3, 4, 6)$ is zero.

More generally, the position $(n_1, \dots, n_k)$ is equivalent to $(n_1 \oplus \dots \oplus n_k)$, since the nim-sum of $(n_1, \dots, n_k, n_1 \oplus \dots \oplus n_k)$ is zero.

Lemma 1 of the previous subsection clearly generalizes to the sum of combinatorial games:

$$(G_1, x_1) \in \mathbf{P} \text{ and } (G_2, x_2) \in \mathbf{N} \text{ imply } (G_1 + G_2, (x_1, x_2)) \in \mathbf{N},$$

$$(G_1, x_1), (G_2, x_2) \in \mathbf{P} \text{ imply } (G_1 + G_2, (x_1, x_2)) \in \mathbf{P}.$$

We also saw that the information $(G_i, x_i) \in \mathbf{N}$ is not enough to decide what kind of position (x_1, x_2) is. Therefore, if we want solve games by dividing them into a sum of smaller games, we need a finer description of the positions than just being in $\mathbf{P}$ or $\mathbf{N}$.

Definition 5 Let G be a progressively bounded combinatorial game in normal play. Its **Sprague-Grundy function** g is defined as follows: for terminal positions t , let $g(t) = 0$, while for other positions,

$$g(x) = \text{mex}\{g(y) : x \rightarrow y \text{ is a legal move}\},$$

where $\text{mex}(S) = \min\{n \geq 0 : n \notin S\}$, for a finite set $S \subseteq \{0, 1, \dots\}$. (This is short for ‘minimal excluded value’).

Note that $g(x) = 0$ is equivalent to $x \in \mathbf{P}$. And a very simple example is that the Sprague-Grundy value of the nim pile (n) is just n .

Theorem 5 (Sprague-Grundy theorem) Every progressively bounded combinatorial game G in normal play is equivalent to a single nim pile, of size $g(x) \geq 0$, where g is the Sprague-Grundy function of G .

We illustrate the theorem with an **example**: a game where a position consists of a pile of chips, and a legal move is to remove 1, 2 or 3 chips. The following table shows the first few values of the Sprague-Grundy function for this game:

x	0	1	2	3	4	5	6
g(x)	0	1	2	3	0	1	2

That is, $g(2) = \text{mex}\{0, 1\} = 2$, $g(3) = \text{mex}\{0, 1, 2\} = 3$, and $g(4) = \text{mex}\{1, 2, 3\} = 0$. In general for this example, $g(x) = x \bmod 4$. We have $(0) \in P_{\text{nim}}$ and $(1), (2), (3) \in N_{\text{nim}}$, hence the $\mathbf{P}$ -positions for our game are the naturals that are divisible by four.

Example: a game consisting of a pile of chips. A legal move from a position with n chips is to remove any positive number of chips strictly smaller than $n/2 + 1$. Here, the first few values of the Sprague-Grundy function are:

x	0	1	2	3	4	5	6
g(x)	0	1	0	2	1	3	0

Definition 6 *The subtraction game with subtraction set $\{a_1, \dots, a_m\}$ is the game in which a position consists of a pile of chips, and a legal move is to remove from the pile a_i chips, for some $i \in \{1, \dots, m\}$.*

The Sprague-Grundy theorem is a consequence of the Sum Theorem just below, by the following simple argument. We need to show that the sum of (G, x) and the single nim pile $(g(x))$ is a **P**-position. By the Sum Theorem and the remarks following Definition 5, the Sprague-Grundy value of this game is $g(x) \oplus g(x) = 0$, which means that is in **P**.

Theorem 6 (Sum Theorem) *If (G_1, x_1) and (G_2, x_2) are two pairs of games and initial starting positions within those games, then, for the sum game $G = G_1 + G_2$, we have that*

$$g(x_1, x_2) = g_1(x_1) \oplus g_2(x_2),$$

where g, g_1, g_2 respectively denote the Sprague-Grundy functions for the games G, G_1 and G_2 .

Proof. First of all, note that if both G_i are progressively bounded, then G is such, too. Hence, we define $B(x_1, x_2)$ to be the maximum number of moves in which the game $(G, (x_1, x_2))$ will end. Note that this quantity is not merely an upper bound on the number of moves, it is the maximum. We will prove the statement by an induction on $B(x_1, x_2) = B(x_1) + B(x_2)$. Specifically, the inductive hypothesis at $n \in \mathbb{N}$ asserts that, for positions (x_1, x_2) in G for which $B(x_1, x_2) \leq n$,

$$g(x_1, x_2) = g_1(x_1) \oplus g_2(x_2). \quad (2)$$

If at least one of x_1 and x_2 is terminal, then (2) is clear: indeed, if x_1 is terminal and x_2 is not, then the game G may only be played in the second coordinate, so it is just the game G_2 in disguise. Suppose then that neither of the positions x_1 and x_2 are terminal ones. We write in binary form:

$$\begin{aligned} g_1(x_1) &= n_1 = n_1^{(m)} n_1^{(m-1)} \dots n_1^{(0)} \\ g_2(x_2) &= n_2 = n_2^{(m)} n_2^{(m-1)} \dots n_2^{(0)}, \end{aligned}$$

so that, for example, $n_1 = \sum_{j=0}^m n_1^{(j)} 2^j$. We know that

$$\begin{aligned} g(x_1, x_2) &= \text{mex}\{g(y_1, y_2) : (x_1, x_2) \rightarrow (y_1, y_2) \text{ a legal move in } G\} \\ &= \text{mex}(A), \end{aligned}$$

where $A := \{g_1(y_1) \oplus g_2(y_2) : (x_1, x_2) \rightarrow (y_1, y_2) \text{ is a legal move in } G\}$. The second equality here follows from the inductive hypothesis, because we know that $B(y_1, y_2) < B(x_1, x_2)$ (the maximum number of moves left in the game G must fall with each move). Writing $s = n_1 \oplus n_2$, we must show that

- (a): $s \notin A$;
- (b): $t \in \mathbb{N}, 0 \leq t < s$ implies that $t \in A$,

since these two statements will imply that $\text{mex}(A) = s$, which yields (2).

Deriving (a): If $(x_1, x_2) \rightarrow (y_1, y_2)$ is a legal move in G , then either $y_1 = x_1$ and $x_2 \rightarrow y_2$ is a legal move in G_2 , or $y_2 = x_2$ and $x_1 \rightarrow y_1$ is a legal move in G_1 . Assuming the first case, we have that

$$g_1(y_1) \oplus g_2(y_2) = g_1(x_1) \oplus g_2(y_2) \neq g_1(x_1) \oplus g_2(x_2),$$

for otherwise, $g_2(y_2) = g_1(x_1) \oplus g_1(x_1) \oplus g_2(y_2) = g_1(x_1) \oplus g_1(x_1) \oplus g_2(x_2) = g_2(x_2)$. This however is impossible, by the definition of the Sprague-Grundy function g_2 , hence $s \notin A$.

Deriving (b): We take $t < s$, and observe that if $t^{(\ell)}$ is the leftmost digit of t that differs from the corresponding one of s , then $t^{(\ell)} = 0$ and $s^{(\ell)} = 1$. Since $s^{(\ell)} = n_1^{(\ell)} + n_2^{(\ell)} \pmod{2}$, we may suppose that $n_1^{(\ell)} = 1$. We want to move in G_1 from x_1 , for which $g_1(x_1) = n_1$, to a position y_1 for which

$$g_1(y_1) = n_1 \oplus s \oplus t. \quad (3)$$

Then we will have $(x_1, x_2) \rightarrow (y_1, x_2)$ on one hand, while

$$g_1(y_1) \oplus g_2(x_2) = n_1 \oplus s \oplus t \oplus n_2 = n_1 \oplus n_2 \oplus s \oplus t = s \oplus s \oplus t = t$$

on the other, hence $t = g_1(y_1) \oplus g_2(x_2) \in A$, as we sought. But why is (3) possible? Well, note that

$$n_1 \oplus s \oplus t < n_1. \quad (4)$$

Indeed, the leftmost digit at which $n_1 \oplus s \oplus t$ differs from n_1 is ℓ , at which the latter has a 1. Since a number whose binary expansion contains a 1 in place ℓ exceeds any number whose expansion has no ones in places ℓ or higher, we see that (4) is valid. The definition of $g_1(x_1)$ now implies that there exists a legal move from x_1 to some y_1 with $g(y_1) = n_1 \oplus s \oplus t$. This finishes case (b) and the proof of the theorem. $\square$

Example. Let G_1 be the subtraction game with subtraction set $S_1 = \{1, 3, 4\}$, G_2 be the subtraction game with $S_2 = \{2, 4, 6\}$, and G_3 be the subtraction game with $S_3 = \{1, 2, \dots, 20\}$. Who has a winning strategy from the starting position $(100, 100, 100)$ in $G_1 + G_2 + G_3$?

2.4 Staircase nim and other examples

Staircase nim. A staircase of n steps contains coins on some of the steps. Let $(x_1, x_2, \dots, x_n)$ denote the position in which there are x_j coins on step j , $j = 1, \dots, n$. A move of staircase nim consists of moving any positive number of coins from any step j to the next lower step, $j - 1$. Coins reaching the ground (step 0) are removed from play. The game ends when all coins are on the ground. Players alternate moves and the last to move wins.

We claim that a configuration is a **P**-position in staircase nim if the numbers of coins on odd-numbered steps forms a **P**-position in nim. To see this, note that moving coins from an odd-numbered step to an even-numbered one represents a legal move in a game of nim consisting of piles of chips lying on the odd-numbered steps. We need only check that moving chips from even to odd numbered steps is not useful. A player who has just seen his opponent to do this may move the chips newly arrived at an odd-numbered location to the next even-numbered one, that is, he may repeat his opponent's move at one step lower. This restores the nim-sum on the odd-numbered steps to its value before the opponent's last move. This means that the extra moves can play no role in changing the outcome of the game from that of nim on the odd-numbered steps.

Moore's nim_k: In this game, recall that players are allowed to remove any number of chips from at most k piles in any given turn. We write the binary expansions of the pile sizes $(n_1, \dots, n_\ell)$:

$$\begin{aligned} n_1 &= n_1^{(m)} \dots n_1^{(0)} \equiv \sum_{j=0}^m n_1^{(j)} 2^j, \\ &\dots \\ n_\ell &= n_\ell^{(m)} \dots n_\ell^{(0)} \equiv \sum_{j=0}^m n_\ell^{(j)} 2^j. \end{aligned}$$

We set

$$\hat{P} = \left\{ (n_1, \dots, n_\ell) : \sum_{i=1}^{\ell} n_i^{(r)} = 0 \pmod{(k+1)} \text{ for each } r \geq 0 \right\}.$$

Theorem 7 (Moore's theorem) *We have $\hat{P} = \mathbf{P}$.*

Proof. Firstly, note that the terminal position 0 lies in $\hat{P}$. There are two other things to check: firstly, that from $\hat{P}$, any legal move takes us out of there. To see this, take any move from a position in $\hat{P}$, and consider the leftmost column for which this move changes the binary expansion of one of the pile numbers. Any change in this column must be from one to zero. The existing sum of the ones and zeros mod $(k+1)$ is zero, and we are adjusting at most k piles. Since ones are turning into zeros, and at least one of them

is changing, we could get back to $0 \bmod k+1$ in this column only if we were to change $k+1$ piles. This isn't allowed, so we have verified that no move from $\hat{P}$ takes us back there.

We must also check that for each position in $\hat{N}$ (which we define to be the complement of $\hat{P}$), there exists a move into $\hat{P}$. This step of the proof is a bit harder. How to select the k piles from which to remove chips? Well, we work by finding the leftmost column whose $\bmod (k+1)$ sum is not-zero. We select any r rows with a one in this column, where r is the number of ones in the column reduced $\bmod (k+1)$ (so that $r \in \{0, \dots, k\}$). We've got the choice to select $k-r$ more rows if we need to. We do this moving to the next column to the right, and computing the number s of ones in that column, ignoring any ones in the rows that we selected before, and reduced $\bmod (k+1)$. If $r+s < k$, then we add s rows to the list of those selected, choosing these so that there is a one in the column currently under consideration, and different from the rows previously selected. If $r+s \geq k$, we choose $k-r$ such rows, so that we have a complete set of k chosen rows. In the first case, we still need more rows, and we collect them successively by examining each successive column to the right in turn, using the same rule as the one we just explained. The point of doing this is that we have chosen the rows in such a way that, for any column, either that column has no ones from the unselected rows because in each of these rows, the most significant digit occurs in a place to the right of this column, or the $\bmod (k+1)$ sum in the rows other than the selected ones is not zero. If a column is of the first type, we set all the bits to zero in the selected rows. This gives us complete freedom to choose the bits in the less significant places. In the other columns, we may have say $t \in \{1, \dots, k\}$ as the $\bmod (k+1)$ sum of the other rows, so we choose the number of ones in the selected rows for this column to be equal to $k-t$. This gives us a $\bmod (k+1)$ sum zero in each row, and thus a position in $\hat{P}$. This argument is not all that straightforward, it may help to try it out on some particular examples: choose a small value of k , make up some pile sizes that lie in $\hat{N}$, and use it to find a specific move to a position in $\hat{P}$. Anyway, that's what had to be checked, and the proof is finished. $\square$

The game of chomp and its solution: A rectangular array of chocolate is to be eaten by two players who alternatively must remove some part of it. A legal move is to choose a vertex and remove that part of the remaining chocolate that lies to the right or above the chosen point. The part removed must be non-empty. The square of chocolate located in the lower-left corner is poisonous, making the aim of the game to force the other player to make the last move. The game is progressively bounded, so that each position is in $\mathbf{N}$ or $\mathbf{P}$. We will show that each rectangular position is in $\mathbf{N}$.

Suppose, on the contrary, that there is a rectangular position in $\mathbf{P}$. Consider the move by player I of chomping the upper-right hand corner. The

resulting position must be in $\mathbf{N}$. This means that player II has a move to $\mathbf{P}$. However, player I can play this move to start with, because each move after the upper-right square of chocolate is gone is available when it was still there. So player I can move to $\mathbf{P}$, a contradiction.

Note that it may not be that chomping the upper-right hand corner is a winning move. This strategy-stealing argument, just as in the case of hex, proves that player I has a winning strategy, without identifying it.

2.5 The game of Green Hackenbush

In the game of Green Hackenbush, we are given a finite graph, that consists of vertices and some undirected edges between some pairs of the vertices. One of the vertices is called the root, and might be thought of as the ground on which the rest of the structure is standing. We talk of ‘green’ Hackenbush because there is an partisan variant of the game in which edges may be colored red or blue instead.

The aim of the players I and II is to remove the last edge from the graph. At any given turn, a player may remove some edge from the graph. This causes not only that edge to disappear, but also all those edges for which every path to the root travels through the edge the player removes.


Note firstly that, if the original graph consists of a finite number of paths, each of which ends at the root, then, in this case, Green Hackenbush is equivalent to the game of nim, where the number of piles is equal to the number of paths, and the number of chips in a pile is equal to the length of the corresponding path.

We need a lemma to handle the case where the graph is a tree:

Lemma 2 (Colon Principle) *The Sprague-Grundy function of Green Hackenbush on a tree is unaffected by the following operation: for any example of two branches of the tree meeting at a vertex, we may replace these two branches by a path emanating from the vertex whose length is the nim-sum of the Sprague-Grundy functions of the two branches.*

Proof. See Ferguson, I-42. The proof in outline: if the two branches consist simply of paths (or ‘stalks’) emanating from a given vertex, then the result is true, by noting that the two branches form a two-pile game of nim, and using the direct sum Theorem for the Sprague-Grundy functions of two games. More generally, we show that we may perform the replacement operation on any two branches meeting at a vertex, by iterating replacing pairs of stalks meeting inside a given branch, until each of the two branches itself has become a stalk. $\square$

As a simple illustration, see the figure. The two branches in this case are stalks, of length 2 and 3. The Sprague-Grundy values of these stalks equal



2 and 3, and their nim-sum is equal to 1. Hence, the replacement operation takes the form shown.

For further discussion of Hackenbush, and references about the game, see Ferguson, Part I, Section 6.

2.6 Wythoff's nim

A position in **Wythoff's nim** consists of a pair of (n, m) of natural numbers, $n, m \geq 0$. A legal move is one of the following: to reduce n to some value between 0 and $n - 1$ without changing m , to reduce m to some value between 0 and $m - 1$ without changing n , or to reduce each of n and m by the same amount, so that the outcome is a pair of natural numbers. The one who reaches $(0, 0)$ is the winner.

Consider the following recursive definition of a sequence of pairs of natural numbers: $(a_0, b_0) = (0, 0)$, $(a_1, b_1) = (1, 2)$, and, for each $k \geq 1$,

$$a_k = \text{mex}\{a_0, a_1, \dots, a_{k-1}, b_0, b_1, \dots, b_{k-1}\}$$

and $b_k = a_k + k$. Each natural number greater than zero is equal to precisely one of the a_i or the b_i . To see this, note that a_j cannot be equal to any of $a_0, \dots, a_{j-1}$ or $b_0, \dots, b_{j-1}$, moreover, for $k > j$ we have $a_k > a_j$ because otherwise a_j would have taken the slot that a_k did. Furthermore, $b_k = a_k + k > a_j + j = b_j$.

It is easy to see that the set of **P** positions is exactly $\{(0, 0), (a_k, b_k), (b_k, a_k), k = 1, 2, \dots\}$. But is there a fast, non-recursive, method to decide if a given position is in **P**?

There is a nice way to construct partitions of the positive integers: fix any irrational $\theta \in (0, 1)$, and set

$$\alpha_k(\theta) = \lfloor k/\theta \rfloor, \quad \beta_k(\theta) = \lfloor k/(1 - \theta) \rfloor.$$

(For rational θ , this definition has to be slightly modified.) Why is this a partition of $\mathbb{Z}_+$? Clearly, $\alpha_k < \alpha_{k+1}$ and $\beta_k < \beta_{k+1}$ for any k . Furthermore, it is impossible to have $k/\theta, \ell/(1 - \theta) \in [N, N + 1)$ for integers k, ℓ, N , because that would easily imply that there are integers in both intervals $I_N = [N\theta, (N + 1)\theta)$ and $J_N = [(N + 1)\theta - 1, N\theta)$, which cannot happen with $\theta \in (0, 1)$. These show that there is no repetition in the set $S = \{\alpha_k, \beta_k, k = 1, 2, \dots\}$. On the other hand, it cannot be that neither of the intervals I_N and J_N contains any integer, and this easily implies $N \in S$, for any N .

Now, we have the question: does there exist a $\theta \in (0, 1)$ for which

$$\alpha_k(\theta) = a_k \quad \text{and} \quad \beta_k(\theta) = b_k? \quad (5)$$

We are going to show that there is only one θ for which this might be true. Since $b_k = a_k + k$, (5) implies that $\lfloor k/\theta \rfloor + k = \lfloor k/(1 - \theta) \rfloor$. Dividing by k and noting that

$$0 \leq k/\theta - \lfloor k/\theta \rfloor < 1,$$

so that

$$0 \leq 1/\theta - (1/k)\lfloor k/\theta \rfloor < 1/k,$$

we find that

$$1/\theta + 1 = 1/(1 - \theta). \quad (6)$$

Thus, $\theta^2 + \theta - 1 = 0$, so that θ or $1/\theta$ equal $2/(1 + \sqrt{5})$. Thus, if there is a solution in $(0, 1)$, it must be this value.

We now define $\theta = 2/(1 + \sqrt{5})$. Note that (6) implies that

$$1/\theta + 1 = 1/(1 - \theta),$$

so that

$$\lfloor k/(1 - \theta) \rfloor = \lfloor k/\theta \rfloor + k.$$

This means that $\beta_k = \alpha_k + k$. We need to verify that

$$\alpha_k = \text{mex}\{\alpha_0, \dots, \alpha_{k-1}, \beta_0, \dots, \beta_{k-1}\}.$$

We checked earlier that α_k is not one of these values. Why is it equal to their mex? Define z to be this mex. If $z \neq \alpha_k$, then $Z < \alpha_k \leq \alpha_l \leq \beta_l$ for all $l \geq k$. Since z is defined as a mex, $z \neq \alpha_i, \beta_i$ for $i \in \{0, \dots, k-1\}$.

3 Two-person zero-sum games

We now turn to studying a class of games that involve two players, with the loss of one equalling the gain of the other in each possible outcome.

3.1 Some examples

A betting game. Suppose that there are two players, a hider and a chooser. The hider has two coins. At the beginning of any given turn, he decides either to place one coin in his left hand, or two coins in his right. He does so, unseen by the chooser, although the chooser is aware that this is the choice that the hider had to make. The chooser then selects one of his hands, and wins the coins hidden there. That means she may get nothing (if the hand is empty), or one or two coins. How should each of the agents play if she wants to maximize her gain, or minimize his loss? Calling the chooser player I and the hider player II , we record the outcomes in a *normal* or *strategic* form:

	II	L	R
I			
L		2	0
R		0	1

If the players choose non-random strategies, and he seeks to minimize his worst loss, while she wants to assure some gain, what are these amounts? In general, consider a pay-off matrix $(a_{i,j})_{i=1,j=1}^{m,n}$, so that player I may play one of m possible plays, and player II one of n possibilities. The meaning of the entries is that a_{ij} is the amount that II pays I in the event that I plays i and II plays j . Let's calculate the assured payment for player I if pure strategies are used. If she announces to player II that she will play i , then II will play that j for which $\min_j a_{ij}$ is attained. Therefore, if she were announcing her choice beforehand, player I would play that i attaining $\max_i \min_j a_{ij}$. On the other hand, if player II has to announce his intention for the coming round to player I , then a similar argument shows that he plays j , where j attains $\min_j \max_i a_{ij}$.

In the example, the assured value for II is 1, and the assured value for I is zero. In plain words, the hider can assure losing only one unit, by placing one coin in his left hand, whereas the chooser knows that he will never lose anything by playing.

It is always true that the assured values satisfy the inequality


$$\min_j \max_i a_{ij} \geq \max_i \min_j a_{ij}.$$

Intuitively, this is because player I cannot be assured of winning more than player II can be guaranteed not to lose. Mathematically, let j^* denote the

value of j that attains the minimum of $\max_i a_{ij}$, and let $\hat{i}$ denote the value of i that attains the maximum of $\min_j a_{ij}$. Then

$$\min_j \max_i a_{ij} = \max_i a_{i\hat{j}} \geq a_{\hat{i}\hat{j}} \geq \min_j a_{\hat{i}j} = \max_i \min_j a_{ij}.$$

If the assured values are not equal, then it makes sense to consider random strategies for the players. Back to our example, suppose that I plays L with probability x and R the rest of the time, whereas II plays L with probability t , and R with probability $1 - t$. Suppose that I announces to II her choice for x . How would II react? If he plays L , his expected loss is $2x$, if R , then $1 - x$. He minimizes the payout and achieves $\min\{2x, 1 - x\}$. Knowing that II will react in this way to hearing the value of x , I will seek to maximize her payoff by choosing x to maximize $\min\{2x, 1 - x\}$. She is choosing the value of x at which the two lines in this graph cross:




So her choice is $x = 1/3$, with which she can assure a payoff $2/3$ on the average, a significant improvement from 0. Looking at things the other way round, suppose that player II has to announce t first. The payoff for player I becomes $2t$ if she picks left and $1 - t$ if she picks right. Player II should choose $t = 1/3$ to minimize his expected payout. This assures him of not paying more than $2/3$ on the average. The two assured values now agree.


Let's look at **another example**. Suppose we are dealing with a game that has the following payoff matrix:

	II	
	L	R
I		
T	0	2
B	5	1

Suppose that player I plays T with probability x and B with probability $1 - x$, and that player II plays L with probability y and R with probability $1 - y$. If player II has declared the value of y , then Player I has expected payoff of $2(1 - y)$ if he plays T , and $4y + 1$ if he plays B . The maximum of these quantities is the expected payoff for player I under his optimal strategy, given that he knows y . Player II minimizes this, and so chooses $y = 1/6$ to obtain an expected payoff of $5/3$.



If player *I* has declared the value of x , then player *II* has expected payment of $5(1 - x)$ if he plays *L* and $1 + x$ if he plays *R*. He minimizes this, and then player *II* chooses x to maximize the resulting quantity. He therefore picks $x = 2/3$, with expected outcome of $5/3$.



In general, player *I* can choose a probability vector

$$x = (x_1, \dots, x_m)^T, \quad \sum_{i=1}^m x_i = 1,$$

where x_i is the probability that he plays i . Player *II* similarly chooses a strategy $y = (y_1, \dots, y_n)^T$. Such randomized strategies are called **mixed**. The resulting expected payoff is given by $\sum x_i a_{ij} y_j = x^T A y$. We will prove **von Neumann's minimax theorem**, which states that

$$\min_y \max_x x^T A y = \max_x \min_y x^T A y.$$

The joint value of the two sides is called the **value** of the game; this is the expected payoff that both players can assure.

3.2 The technique of domination

We illustrate a useful technique with another example. Two players choose numbers in $\{1, 2, \dots, n\}$. The player whose number is higher than that of her opponent by one wins a dollar, but if it exceeds the other number by two or more, she loses 2 dollars. In the event of a tie, no money changes hands. We write the payoff matrix for the game:

II	1	2	3	4	$\dots$	n
I						
1	0	-1	2	2	$\dots$	2
2	1	0	-1	2	$\dots$	2
3	-2	1	0	-1	2	$\dots$ 2
$\cdot$						
$\cdot$						
$n-1$	-2	-2	$\dots$		1	0 -1
n	-2	-2	$\dots$			1 0

This apparently daunting example can be reduced by a new technique, that of domination: if row i has each of its elements at least the corresponding element in row $\hat{i}$, that is, if $a_{ij} \geq a_{\hat{i}j}$ for each j , then, for the purpose of determining the value of the game, we may erase row $\hat{i}$. (The value of the game is defined as the value arising from von Neumann's minimax theorem). Similarly, there is a notion of domination for player II : if $a_{ij} \leq a_{ij^*}$ for each i , then we can eliminate column j^* without affecting the value of the game.

Let us see in details why this is true. Assuming that $a_{ij} \leq a_{ij^*}$ for each i , if player II changes a mixed strategy y to another z by letting $z_j = y_j + y_{j^*}$, $z_{j^*} = 0$ and $z_\ell = y_\ell$ for all $\ell \neq j, j^*$, then

$$\sum_{i,\ell} x_i a_{i,\ell} y_\ell = x^T A y \geq \sum_{i,\ell} x_i a_{i,\ell} z_\ell = x^T A z,$$

because $\sum_i x_i (a_{i,j} y_j + a_{i,j^*} y_{j^*}) \geq \sum_i x_i a_{i,j} (y_j + y_{j^*})$. Therefore, strategy z , in which she didn't use column j^* , is at least as good for player II as y .

In the example in question, we may eliminate each row and column indexed by four or greater. We obtain the reduced game:

<i>II</i>	1	2	3
<i>I</i>			
1	0	-1	2
2	1	0	-1
3	-2	1	0

Consider (x_1, x_2, x_3) as a strategy for player I . The expected payments made by player II under her pure strategies 1, 2 and 3 are

$$(x_2 - 2x_3, -x_1 + x_3, 2x_1 - x_3). \quad (7)$$

Player II seeks to minimize her expected payment. Player I is choosing (x_1, x_2, x_3) : for the time being, suppose that she fixes x_3 , and optimizes her choice for x_1 . Eliminating x_2 , (7) becomes

$$(1 - x_1 - 3x_3, -x_1 + x_3, 3x_1 + x_3 - 1).$$

Computing the choice of x_1 for which the maximum of the minimum of these quantities is attained, and then maximising this over x_3 , yields an optimal strategy for each player of $(1/4, 1/2, 1/4)$, and a value for the game of zero.

Remark. It can of course happen in a game that none of the rows dominates another one, but there are two rows, v, w , whose convex combination $pv + (1 - p)w$ for some $p \in (0, 1)$ does dominate some other rows. In this case the dominated rows can still be eliminated.

3.3 The use of symmetry

We illustrate a symmetry argument by analysing **the game of battleship and salvo**:

X	X	
	B	

A battleship is located on two adjacent squares of a three-by-three grid, shown by the two X s in the example. A bomber, who cannot see the submerged craft, hovers overhead. He drops a bomb, denoted by B in the figure, on one of the nine squares. He wins if he hits and loses if he misses the submarine. There are nine pure strategies for the bomber, and twelve for the submarine. That means that the payoff matrix for the game is pretty big. We can use symmetry arguments to simplify the analysis of the game.

Indeed, suppose that we have two bijections

$$g_1 : \{\text{moves of } I\} \rightarrow \{\text{moves of } I\}$$

and

$$g_2 : \{\text{moves of } II\} \rightarrow \{\text{moves of } II\},$$

for which the payoffs a_{ij} satisfy

$$a_{g_1(i), g_2(j)} = a_{ij}. \quad (8)$$

If this is so, then there are optimal strategies for player I that give equal weight to $g_1(i)$ and i for each i . Similarly, there exists a mixed strategy for player II that is optimal and assigns the same weight to the moves $g_2(j)$ and j for each j .

In the example, we may take g_1 to be the map that flips the first and the third columns. Similarly, we take g_2 to do this, but for the battleship location. Another example of a pair of maps satisfying (8) for this game: g_1 rotates the bomber's location by 90 degrees anticlockwise, whereas g_2 does the same for the location of the battleship. Using these two symmetries, we may now write down a much more manageable payoff matrix:

SHIP BOMBER	center	off-center
corner	0	1/4
midside	1/4	1/4
middle	1	0

These are the payoff in the various cases for play of each of the agents. Note that the pure strategy of corner for the bomber in this reduced game in fact corresponds to the mixed strategy of bombing each corner with $1/4$ probability in the original game. We have a similar situation for each of the pure strategies in the reduced game.

We use domination to simplify things further. For the bomber, the strategy ‘midside’ dominates that of ‘corner’. We have busted down to:

SHIP BOMBER	center	off-center
midside	1/4	1/4
middle	1	0

Now note that for the ship (that is trying to escape the bomb and thus is heading away from the high numbers on the table), off-center dominates center, and thus we have the reduced table:

SHIP BOMBER	off-center
midside	1/4
middle	0

The bomber picks the better alternative — technically, another application of domination — and picks midside over middle. The value of the game is $1/4$, the bomb drops on one of the four middles of the sides with probability $1/4$ for each, and the submarine hides in one of the eight possible locations that exclude the center, choosing any given one with a probability of $1/8$.

3.4 von Neumann’s minimax theorem

We begin this section with some preliminaries of the proof of the minimax theorem. We mentioned that convex geometry plays an important role in the von Neumann minimax theorem. Recall that:

Definition 7 A set $K \subseteq \mathbb{R}^d$ is convex if, for any two points $\mathbf{a}, \mathbf{b} \in K$, the line segment that connects them,

$$\{p\mathbf{a} + (1 - p)\mathbf{b} : p \in [0, 1]\},$$

also lies in K .

The main fact about convex sets that we will need is:

Theorem 8 (Separation theorem for convex sets) *Suppose that $K \subseteq \mathbb{R}^d$ is closed and convex. If $\mathbf{0} \notin K$, then there exists $\mathbf{z} \in \mathbb{R}^d$ and $c \in \mathbb{R}$ such that*

$$0 < c < \mathbf{z}^T \mathbf{v},$$

for all $\mathbf{v} \in K$.

What the theorem is saying is that there is a hyperplane that separates $\mathbf{0}$ from K : this means a line in the plane, or a plane in $\mathbb{R}^3$. The hyperplane is given by

$$\{\mathbf{x} \in \mathbb{R}^d : \mathbf{z}^T \mathbf{x} = c\}.$$

The theorem implies that on any continuous path from $\mathbf{0}$ to K , there is some point that is on this hyperplane.

Proof: There is $\mathbf{z} \in K$ for which

$$\|\mathbf{z}\| = \sqrt{\sum_{i=1}^d z_i^2} = \inf_{\mathbf{v} \in K} \|\mathbf{v}\|.$$

This is because the function $\mathbf{v} \mapsto \|\mathbf{v}\|$, considered as $K \cap \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq R\} \rightarrow [0, \infty)$, is continuous, with its domain being a closed and bounded set, which is non-empty if R is large enough. Therefore, the map attains its infimum, at a point that we have called $\mathbf{z}$. Since $\|\mathbf{z}\| \leq R$, there can be no point with a lower norm that is in the part of K not in the domain of this map.

Now choose $c = (1/2)\|\mathbf{z}\|^2 > 0$. We have to check that $c < \mathbf{z}^T \mathbf{v}$ for each $\mathbf{v} \in K$. To do so, consider such a $\mathbf{v}$. For $\epsilon \in (0, 1)$, we have that $\epsilon \mathbf{v} + (1 - \epsilon)\mathbf{z} \in K$, because $\mathbf{z}, \mathbf{v} \in K$ and K is convex. Hence,

$$\|\mathbf{z}\|^2 \leq \|\epsilon \mathbf{v} + (1 - \epsilon)\mathbf{z}\|^2 = (\epsilon \mathbf{v}^T + (1 - \epsilon)\mathbf{z}^T)(\epsilon \mathbf{v} + (1 - \epsilon)\mathbf{z}),$$

the first inequality following from the fact that $\mathbf{z}$ has the minimum norm of any point in K . We obtain

$$\mathbf{z}^T \mathbf{z} \leq \epsilon^2 \mathbf{v}^T \mathbf{v} + (1 - \epsilon)^2 \mathbf{z}^T \mathbf{z} + 2\epsilon(1 - \epsilon)\mathbf{v}^T \mathbf{z}.$$

Multiplying out and cancelling an ϵ :

$$\epsilon(2\mathbf{v}^T \mathbf{z} - \mathbf{v}^T \mathbf{v} - \mathbf{z}^T \mathbf{z}) \leq 2(\mathbf{v}^T \mathbf{z} - \mathbf{z}^T \mathbf{z}).$$

Taking $\epsilon \downarrow 0$, we find that

$$0 \leq \mathbf{v}^T \mathbf{z} - \mathbf{z}^T \mathbf{z},$$

which implies that

$$\mathbf{v}^T \mathbf{z} \geq 2c > c,$$

as required. $\square$

The minimax theorem shows that two-person zero-sum games have a value, in a sense we now describe. Suppose given a payoff matrix $A = (a_{ij})_{i,j=1}^{m,n}$, with a_{ij} equal to the payment of II to I if I picks i and II picks j . Denote by

$$\Delta_m = \{x \in \mathbb{R}^m : x_i \geq 0, \sum_{i=1}^m x_i = 1\}$$

the set of all probability distributions on m values. Now, player I can assure an expected payoff of $\max_{x \in \Delta_m} \min_{y \in \Delta_n} x^T A y$. On the other hand, player II can be assured of not paying more than $\min_{y \in \Delta_n} \max_{x \in \Delta_m} x^T A y$ on the average. Having prepared some required tools, we will now prove:

Theorem 9 (von Neumann minimax)

$$\max_{x \in \Delta_m} \min_{y \in \Delta_n} x^T A y = \min_{y \in \Delta_n} \max_{x \in \Delta_m} x^T A y.$$

Proof: The direction

$$\max_{x \in \Delta_m} \min_{y \in \Delta_n} x^T A y \leq \min_{y \in \Delta_n} \max_{x \in \Delta_m} x^T A y$$

is easy, and has the same proof as that given for pure strategies. For the other inequality, we firstly assume that

$$\min_{y \in \Delta_n} \max_{x \in \Delta_m} x^T A y > 0. \quad (9)$$

Let K denote the set of payoff vectors that player II can achieve, or that are better for player I than some such vector. That is,

$$\begin{aligned} K = & \left\{ A y + v = y_1 A^{(1)} + y_2 A^{(2)} + \dots + y_n A^{(n)} + v : \right. \\ & \left. y = (y_1, \dots, y_n)^T \in \Delta_n, v = (v_1, \dots, v_m)^T, v_i \geq 0 \right\}, \end{aligned}$$

where $A^{(i)}$ denotes the i th column of A . It is easy to show that K is convex and closed: this uses the fact that Δ_n is closed and bounded. Note also that

$$\mathbf{0} \notin K. \quad (10)$$

To see this, note that (9) means that, for every y , there exists x such that player I has a uniformly positive expected payoff $x^T A y > \delta > 0$. If $\mathbf{0} \in K$, this means that, for some y , we have that

$$A y = y_1 A^{(1)} + y_2 A^{(2)} + \dots + y_n A^{(n)} \leq 0,$$

where by ≤ 0 we mean that in each of the coordinates. However, this contradicts (9), and proves (10).

The separation theorem now allows us to find $z \in \mathbb{R}^m$ and $c > 0$ such that $0 < c < z^T w$ for all $w \in K$. That is, $z^T(Ay + v) > c > 0$ for all $y \in \Delta_n$, $v \geq 0$, $v \in \mathbb{R}^m$. We deduce that $z_i \geq 0$ by considering large and positive choices for v_i , and using $z^T v = z_1 v_1 + \dots + z_m v_m$. Moreover, not all of the z_i can be zero. This means that we may set $s = \sum_{i=1}^m z_i > 0$, so that $x = (1/s)(z_1, \dots, z_m)^T = (1/s)z \in \Delta_m$ with $x^T Ay > 0$ for all $y \in \Delta_n$. Since Δ_n is closed, we also have $\min_{y \in \Delta_n} x^T Ay > 0$ for this $x \in \Delta_m$.

We now remove the assumption (9), and suppose, playing on contradiction, that

$$\max_{x \in \Delta_m} \min_{y \in \Delta_n} x^T Ay < \lambda < \min_{y \in \Delta_n} \max_{x \in \Delta_m} x^T Ay.$$

Consider a new game with payoff matrix $\hat{A}$ given by $\hat{A}_{m \times n} = (\hat{a}_{ij})_{1 \leq i \leq m, 1 \leq j \leq n}$, with $\hat{a}_{ij} = a_{ij} - \lambda$. We find that

$$\max_{x \in \Delta_m} \min_{y \in \Delta_n} x^T \hat{A} y < 0 < \min_{y \in \Delta_n} \max_{x \in \Delta_m} x^T \hat{A} y.$$

The right hand inequality says that $\hat{A}$ satisfies (9), so, as we saw above, we also have $\max_{x \in \Delta_m} \min_{y \in \Delta_n} x^T \hat{A} y > 0$. But this contradicts the left hand inequality. $\square$

The above proof shows only the existence of the minimax, but does not tell how one can actually find it. It can be shown that finding the minimax in general is equivalent to solving a general linear program, which is known to be an algorithmically difficult task.

3.5 Resistor networks and troll games

Suppose that two cities, A and B , are connected by a ‘parallel-series’ network of roads. Such a network is built by modifying an initial straight road that runs from A to B . The modification takes the form of a finite number of steps, these steps being of either *series* or *parallel* type. A series step is to find a road in the current network and to replace it by two roads that run one after the other along the current path of the road. The parallel step consists of replacing one current road by two, each of which runs from the start to the finish of the road being replaced.

Imagine that each road has a cost attached to it. A troll and a traveller will each travel from city A to city B along some route of their choice. The traveller will pay the troll the total cost of the roads that both of them choose.

This is of course a zero-sum game. However, as we shall see in a minute, it is of a quite special type, and there is a nice general way to solve such games, as follows.

We interpret the road network as an electrical one, and the costs as resistances. Resistances add in series, whereas conductances, which are their reciprocals, add in parallel. We claim that the optimal strategies for both players are the same. Under the optimal strategy, a player who reaches a fork should move along any of the edges emanating from the fork towards B with a probability proportional to the conductance of that edge.

A way of seeing why these games are solved like this is to introduce the notion of two games being summed in parallel or in series. Suppose given two zero-sum games G_1 and G_2 with values v_1 and v_2 . Their **series addition** just means: play G_1 , and then G_2 . The series sum game has value $v_1 + v_2$. In the **parallel-sum game**, each player chooses either G_1 or G_2 to play. If each picks the same game, then it is that game which is played. If they differ, then no game is played, and the payoff is zero. We may write a big payoff matrix as follows:

	II	
I		
	G_1	0
	0	G_2

If the two players play G_1 and G_2 optimally, the payoff matrix is effectively:

	II	
I		
	v_1	0
	0	v_2


The optimal strategy for each player consists of playing G_1 with probability $v_2/(v_1 + v_2)$, and G_2 with probability $v_1/(v_1 + v_2)$. Given that

$$\frac{v_1 v_2}{v_1 + v_2} = \frac{1}{1/v_1 + 1/v_2},$$

this explains the form of the optimal strategy in troll-traveller games on series-parallel graphs.

Example: see the figure. It shows that the game on the network in the upper left corner, with resistances all equaling to 1, has value $3/5$.

On general graphs with two distinguished vertices A, B , we need to define the game in the following way: if the troll and the traveller traverse an edge in the opposite directions, then the troll pays the cost of the road to the traveller. Then the value of the game turns out to be the **effective resistance** between A and B , a quantity with important meaning in several probabilistic contexts.



Adding in series and parallel for the troll-traveller game

3.6 Hide-and-seek games

Games of hide and seek form another class of two-person zero-sum games that we will analyse. For this, we need a tool:

Lemma 3 (Hall's marriage lemma) *Suppose given a set B of boys and a set G of girls, with $|B| = b$ and $|G| = g$ satisfying $g \geq b$. Let $f : B \rightarrow 2^G$ be such that $f(c)$ denotes the subset of G that are the girls known to boy c . Each boy can be matched to a girl that he knows if and only if, for each $B' \subseteq B$, we have that $|f(B')| \geq |B'|$. (Here, the function f has been extended to subsets B' of B by setting $f(B') = \cup_{c \in B'} f(c)$.)*

Proof: It is clear from the definition of a matching that the existence of a matching implies the other condition. We will prove the opposite direction by an induction on b , the number of boys. The case when $b = 1$ is easy. For larger values of b , suppose that the statement we seek is known for $b' < b$. Two cases: firstly, suppose that there exists $B' \subseteq B$ satisfying $|f(B')| = |B'|$ with $|B'| < b$. We perform a matching of B' to $f(B')$ by using the inductive hypothesis. If $A \subseteq B \setminus B'$, then $|f(A) \setminus f(B')| \geq |A|$; this is because $|f(A \cup B')| = |f(B')| + |f(A) \setminus f(B')|$ and $|f(A \cup B')| \geq |A \cup B'| = |A| + |B'|$ by assumption. Hence, we may apply the inductive hypothesis to the set $B \setminus B'$ to find a matching of this set to girls in $G \setminus f(B')$. We have found a matching of B into G as required.

In the second case, $|f(B')| > |B'|$ for each $B' \subseteq B$. This case is easy: we just match a given boy to any girl he knows. Then the set of remaining boys still satisfies the second condition in the statement of the lemma. By the inductive hypothesis, we match them, and we have finished the proof. $\square$

We now prove a valuable result, a version of **König's lemma**, by using Hall's marriage lemma.

Lemma 4 *Consider an $n \times m$ matrix whose entries consist of 0s and 1s. Call two 1s independent if no row nor column contain them both. A cover of the matrix is a collection of rows and column whose union contains each of the 1s. Then: the maximal size of a set of independent 1s is equal to the minimal size of a cover.*

Proof: Consider a maximal independent set of 1s (of size k), and a minimal cover consisting of ℓ lines. That $k \leq \ell$ is easy: each 1 in the independent set is covered by a line, and no two are covered by the same line. For the other direction we make use of Hall's lemma. Suppose that among these ℓ lines, there are r rows and c columns. In applying Hall's lemma, the rows correspond to the boys and columns not in the cover to girls. A row knows such a column if their intersection contains a 1 from the independent set.

Suppose that j of these rows know $s < j$ columns not in the minimal cover. We could replace these j rows by these s columns to obtain a smaller cover. This is impossible, meaning that every set of j rows has to know at least j columns not in the minimal cover. By Hall's lemma, we can match up the r rows with columns outside the cover and known to them.

Similarly, we obtain a 1 – 1 matching of the c columns in the cover with c rows outside the cover. Each of the intersections of these c matched rows and columns contains a 1. Similarly, with the r matched rows and columns just constructed. The $r + c$ resulting 1s are independent, hence $k \geq \ell$. This completes the proof. $\square$

We use König's lemma to analyse a **hide-and-seek** game in a matrix. In this game, a matrix whose entries are 0s and 1s is given. Player *I* chooses a 1 somewhere in the matrix, and hides there. Player *II* chooses a row or column and wins a payoff of 1 if the line that he picks contains the location chosen by player *I*. A strategy for player *I* is to pick a maximal independent set of 1s, and then hide in a uniformly chosen element of it. A strategy for player *II* consists of picking uniformly at random one of the lines of a minimal cover of the matrix. König's lemma shows that this is a joint optimal strategy, and that the value of the game is k^{-1} , where k is the size of the maximal set of independent 1s.

3.7 General hide-and-seek games

We now analyse a more general version of the game of hide-and-seek. A matrix of values $(b_{ij})_{n \times n}$ is given. Player *II* chooses a location (i, j) at which to hide. Player *I* chooses a row or a column of the matrix. He wins a payment of b_{ij} if the line he has chosen contains the hiding place of his opponent.

Firstly, we propose a strategy for player *II*, later checking that it is optimal. The player may choose a fixed permutation π of the set $\{1, \dots, n\}$ and then hide at location (i, π_i) with a probability p_i that he chooses. Given

a choice π , the optimal choice for p_i is $p_i = d_{i,\pi_i}/D_\pi$, where $d_{ij} = b_{ij}^{-1}$ and $D_\pi = \sum_{i=1}^n d_{i,\pi_i}$, because it is this choice that equalizes the expected payments. The expected payoff for the game is then $1/D_\pi$.

Thus, if Player *II* is going to use a strategy that consists of picking a permutation π^* and then doing as described, the right permutation to pick is one that maximizes D_π . We will in fact show that doing this is an optimal strategy, not just in the restricted class of those involving permutations in this way, but over all possible strategies.

To find an optimal strategy for Player *I*, we need an analogue of König's lemma. In this context, a *covering* of the matrix $D = (d_{ij})_{n \times n}$ will be a pair of vectors $u = (u_1, \dots, u_n)$ and $w = (w_1, \dots, w_n)$ such that $u_i + w_j \geq d_{ij}$ for each pair (i, j) . (We assume that u and w have non-negative components). The analogue of the König lemma:

Lemma 5 *Consider a minimal covering (u^*, w^*) . (This means one for which $\sum_{i=1}^n (u_i + w_i)$ is minimal). Then:*

$$\sum_{i=1}^n (u_i^* + w_i^*) = \max_{\pi} D_\pi. \quad (11)$$

Proof: Note firstly that a minimal covering exists, because the map

$$(u, w) \mapsto \sum_{i=1}^n (u_i + w_i),$$

defined on the closed and bounded set $\{(u, w) : 0 \leq u_i, w_i \leq M, u_i + w_j \geq d_{ij}\}$, where $M = \max_{i,j} d_{i,j}$, does indeed attain its infimum.

Note also that we may assume that $\min_i u_i^* > 0$.

That the left-hand-side of (11) is at least the right-hand-side is straightforward. Indeed, for any π , we have that $u_i^* + w_{\pi_i}^* \geq d_{i,\pi_i}$. Summing over i , we obtain this inequality.

Showing the other inequality is harder, and requires Hall's marriage lemma, or something similar. We need a definition of 'knowing' to use the Hall lemma. We say that row i knows column j if

$$u_i^* + w_j^* = d_{ij}.$$

Let's check Hall's condition. Suppose that k rows $i_1, \dots, i_k$ know between them only $\ell < k$ columns $j_1, \dots, j_\ell$. Define $\tilde{u}$ from u^* by reducing these rows by a small amount ϵ . Leave the other rows unchanged. The condition that ϵ must satisfy is in fact that

$$\epsilon \leq \min_i u_i^*$$

and also

$$\epsilon \leq \min \{u_i + w_j - d_{ij} : (i, j) \text{ such that } u_i + w_j > d_{ij}\}.$$

Similarly, define $\tilde{w}$ from w^* by adding ϵ to the ℓ columns known by the k rows. Leave the other columns unchanged. That is, for the columns that are changing,

$$\tilde{w}_{j_i} = w_{j_i}^* + \epsilon \text{ for } i \in \{1, \dots, \ell\}.$$

We claim that $(\tilde{u}, \tilde{w})$ is a covering of the matrix. At places where the equality $d_{ij} = u_i^* + w_j^*$ holds, we have that $d_{ij} = \tilde{u}_i + \tilde{w}_j$, by construction. In places where $d_{ij} < u_i^* + w_j^*$, then

$$\tilde{u}_i + \tilde{w}_j \geq u_i^* - \epsilon + w_j^* > d_{ij},$$

the latter inequality by the assumption on the value of ϵ .

The covering $(\tilde{u}, \tilde{w})$ has a strictly smaller sum of components than does (u^*, w^*) , contradicting the fact that this latter covering was chosen to be minimal.

We have checked that Hall's condition holds. Hall's lemma provides a matching of columns and rows. This is a permutation π^* such that, for each i , we have that

$$u_i^* + w_{\pi(i)}^* = d_{i, \pi(i)},$$

from which it follows that

$$\sum_{i=1}^n u_i^* + \sum_{i=1}^n w_i^* = D_{\pi^*}.$$

We have found a permutation π^* that gives the other inequality required to prove the lemma. $\square$

This lemma shows a pair of optimal strategies for the players. Player I chooses row i with probability u_i^*/D_{π^*} , and column j with probability w_j^*/D_{π^*} . Against this strategy, if player II chooses some (i, j) , then the payoff will be

$$\frac{u_i^* + v_j^*}{D_{\pi^*}} b_{ij} \geq \frac{d_{ij} b_{ij}}{D_{\pi^*}} = D_{\pi^*}^{-1}.$$

We deduce that the permutation strategy for player II described before the lemma is indeed optimal.

Example. Consider the hide-and-seek game with payoff matrix B given by

$$\begin{vmatrix} 1 & 1/2 \\ 1/3 & 1/5 \end{vmatrix}$$

This means that the matrix D is equal to

$$\begin{vmatrix} 1 & 2 \\ 3 & 5 \end{vmatrix}$$

To determine a minimal cover of the matrix D , consider first a cover that has all of its mass on the rows: $u = (2, 5)$ and $v = (0, 0)$. Note that rows 1 and 2 know only column 2, according to the definition of ‘knowing’ introduced in the analysis of this game. Modifying the vectors u and v according to the rule given in this analysis, we obtain updated vectors, $u = (1, 4)$ and $v = (0, 1)$, whose sum is 6, equal to the expression $\max_{\pi} D_{\pi}$ (obtained by choosing the permutation $\pi = id$.)


An optimal strategy for the hider is to play $p(1, 1) = 1/6$ and $p(2, 2) = 5/6$. An optimal strategy for the seeker consists of playing $q(row1) = 1/6$, $q(row2) = 2/3$ and $q(col2) = 1/6$. The value of the game is $1/6$.

3.8 The bomber and submarine game

In the bomber and submarine game, a submarine is located at a site in $\mathbb{Z}$ at any given time step in $\{0, 1, \dots\}$, and moves either left or right for the next time step. In the game G_n , the bomber drops one bomb at some time $j \in \{0, 1, \dots, n\}$ at some site in $\mathbb{Z}$. The bomb arrives at time $j + 2$, and destroys the submarine if it hits the right site. What is the value of the game G_n ? The answer depends on n . The value of G_0 is $1/3$, because the submarine may ensure that it has a $1/3$ probability of being at any of the sites $-2, 0$ or 2 at time 2 . It moves left or right with equal probability at the first time step, and then turns with probability of $1/3$. The value of $1/3$ for the game G_1 can be obtained by pursuing the above strategy. We have already decided how the submarine should move in the first two time steps. If the submarine is at 1 at time 1 , then it moves to 2 at time 2 with probability $2/3$. Thus, it should move with probability $1/2$ to each of sites 1 or 3 at time 2 if it is at site 2 at that time, to ensure a value of $1/3$ for G_1 . This forces it to move from site 0 to site 1 with probability 1 , if it visited site 1 at time 1 . Obtaining the same values in the symmetric case where the submarine moves through site -1 at time 1 , we obtain a strategy for the submarine that ensures that it is hit with a maximum probability of $1/3$ in G_1 .

It is impossible to pursue this strategy to obtain a value of $1/3$ for the game G_2 . Indeed, $v(G_2) > 1/3$. We now describe Isaac’s strategy for the game. It is not optimal in any given game G_n , but it does have the merit of having the same limiting value, as $n \rightarrow \infty$, as optimal play. In G_0 , the strategy is as shown:

The rule in general is: turn with a probability of $1 - a$, and keep going with a probability of a . The strategy is simple in the sense that the transition rates in any 2-play subtree of the form in the figure has the transition rates shown there, or its mirror image. We now choose a to optimize the probability of evasion for the submarine. Its probabilities of arrival at sites $-2, 0$ or 2 at time 2 are a^2 , $1 - a$ and $a(1 - a)$. We have to choose a so that $\max\{a^2, 1 - a\}$ is minimal. This value is achieved when $a^2 = 1 - a$, whose



solution in $(0, 1)$ is given by $a = 2/(1 + \sqrt{5})$.

The payoff for the bomber against this strategy is at most $1 - a$. We have proved that the value $v(G_n)$ of the game G_n is at most $1 - a$, for each n .

3.9 A further example

Consider the zero sum game whose payoff matrix is given by:

<i>II</i>			
<i>I</i>			
	1	0	8
	2	3	-1

To solve this game, firstly, we search for saddle points — a value in the matrix that is maximal in its column and minimal in its row. None exist in this case. Nor are there any evident dominations of rows or columns.

Suppose then that player *I* plays the mixed strategy $(p, 1 - p)$. If there is an optimal strategy for player *II* in which she plays each of her three pure strategies with positive probability, then

$$2 - p = 3 - 3p = 9p - 1.$$

No solution exists, so we consider now mixed strategies for player *II* in which one pure strategy is never played. If the third column has no weight, then $2 - p = 3 - 3p$ implies that $p = 1/2$. However, the entry 3 in the matrix becomes a saddle point in the 2×2 matrix formed by eliminating the third column, which is not consistent with $p = 1/2$.

Consider instead strategies supported on columns 1 and 3. The equality $2 - p = 9p - 1$ yields $p = 3/10$, giving a payoff for player *II* of

$$(17/10, 27/10, 17/10).$$

If player *II* plays column 1 with probability x and column 3 otherwise, then player *I* sees the payoff vector $(8 - 7x, 3x - 1)$. These quantities are equal when $x = 9/10$, so that player *I* sees the payoff vector $(17/10, 17/10)$. Thus, the value of the game is $17/10$.

4 General sum games

We now turn to discussing the theory of general sum games. Such a game is given in strategic form by two matrices A and B , whose entries give the payoffs of given joint pure strategies to each of the two players. Usually there is no joint optimal strategy for the players, but still exists a generalization of the von Neumann minimax, the so-called Nash equilibrium. These equilibria give the strategies that “rational” players could follow. However, there are often several Nash equilibria, and in choosing one of them, some degree of cooperation between the players may be optimal. Moreover, a pair of strategies based on cooperation might be better for both players than any of the Nash equilibria. We begin with two examples.

4.1 Some examples

Example: the prisoner’s dilemma. Two suspects are held and questioned by police who ask each of them to confess or to remain silent. The charge is serious, but the evidence held by the police is poor. If one confesses and the other is silent, then the first goes free, and the other is sentenced to ten years. If both confess, they will each spend eight years in prison. If both remain silent, one year to each is the sentence, for some minor crime that the police is able to prove. Writing the payoff as the number of years that remain in the following decade apart from those spent in prison, we obtain the following payoff matrix:

II	S	C
I		
S	(9,9)	(0,10)
C	(10,0)	(2,2)

The pay-off matrices for players I and II are the 2×2 matrices given by the collection of first, or second, entries in each of the vectors in the above matrix.

If the players only play one round, then there is an argument involving domination saying that each should confess: the outcome she secures by confessing is preferable to the alternative of remaining silent, whatever the behaviour of the other player. However, this outcome is much worse for each player than the one achieved by both remaining silent. In a once-only game, the ‘globally’ preferable outcome of each remaining silent could only occur were each player to suppress the desire to achieve the best outcome in selfish terms. In games with repeated play ending at a known time, the same applies, by an argument of backward induction. In games with repeated play ending at a random time, however, the globally preferable solution may arise even with selfish play.

Example: the battle of the sexes. The wife wants to head to the opera but the husband yearns instead to spend an evening watching baseball. Neither is satisfied by an evening without the other. In numbers, player I being the wife and II the husband, here is the scenario:

	II	O	B
I			
O		(4,1)	(0,0)
B		(0,0)	(1,4)

One might naturally come up with two modifications of von Neumann's minimax. The first one is that the players do not suppose any rationality about their partner, so they just want to assure a payoff assuming the worst-case scenario. Player I can guarantee a safety value of $\max_{p \in \Delta_2} \min_{q \in \Delta_2} p^T A q$, where A denotes the matrix of payoffs received by her. This gives the strategy $(1/5, 4/5)$ for her, with an assured payoff $4/5$, which value actually does not depend on what player II does. The analogous strategy for player II is $(4/5, 1/5)$, with the same assured payoff $4/5$. Note that these values are lower than what each player would get from just agreeing to go where the other prefers.

The second possible adaptation of the minimax approach is that player I announces her value p , expecting player II to maximize his payoff given this p . Then player I maximizes the result over p . However, in contrast to the case of zero-sum games, the possibility of announcing a strategy and committing to it in a general-sum game might actually raise the payoff for the announcer, and hence it becomes a question how a model can accommodate this possibility. In our game, each player could just announce their favorite choice, and to expect their spouse to behave "rationally" and agree with them. This leads to a disaster, unless one of them manages to make this announcement before the spouse does, and the spouse truly believes that this decision is impossible to change, and takes the effort to act rationally.

In this example, it is quite artificial to suppose that the two players cannot discuss, and that there are no repeated plays. Nevertheless, this example shows clearly that a minimax approach is not suitable any more.

4.2 Nash equilibrium

We now introduce a central notion for the study of general sum games:

Definition 8 (Nash equilibrium) A pair of vectors (x^*, y^*) with $x^* \in \Delta_m$ and $y^* \in \Delta_n$ define a **Nash equilibrium**, if no player gains by deviating unilaterally from it. That is,

$$x^{*T} A y^* \geq x^T A y^*$$

for all $x \in \Delta_m$, and

$$x^{*T} B y^* \geq x^{*T} B y$$

for all $y \in \Delta_n$. The game is called **symmetric** if $m = n$ and $A_{ij} = B_{ji}$ for all $i, j \in \{1, 2, \dots, n\}$. A pair (x, y) of strategies is called *symmetric* if $x_i = y_i$ for all $i = 1, \dots, n$.

We will see that there always exists a Nash equilibrium; however, there can be many of them. If x and y are unit vectors, with a 1 in some coordinate and 0 in the others, then the equilibrium is called **pure**.

In the above example of the battle of the sexes, there are two pure equilibria: these are BB and OO. There is also a mixed equilibrium, $(4/5, 1/5)$ for player I and $(1/5, 4/5)$ for II, having the value $4/5$, which is very low, again.

Example: lions, cheetahs and antelopes. Antelopes have been observed to jump energetically when a lion nearby seems liable to hunt them. Why do they expend energy in this way? One theory was that the antelopes are signalling danger to others at some distance, in a community-spirited gesture. However, the antelopes have been observed doing this all alone. The currently accepted theory is that the signal is intended for the lion, to indicate that the antelope is in good health and is unlikely to be caught in a chase. Another justification of this explanation is that the antelopes do not jump when cheetahs are approaching, but start running as fast as they can. Being healthy is only a necessary, not a sufficient condition to escape from a cheetah.

Consider a simple model, where two cheetahs are giving chase to two antelopes. The cheetahs will catch any antelope they choose. If they choose the same one, they must share the spoils. If they catch one without the other cheetah doing likewise, the catch is unshared. There is a large antelope and a small one, that are worth ℓ and s to the cheetahs. Here is the matrix of payoffs:

I \ II	L	S
	L	S
L	$(\ell/2, \ell/2)$	(ℓ, s)
S	(s, ℓ)	$(s/2, s/2)$

If $\ell \geq 2s$, the first row dominates the second, and likewise, the columns; hence each cheetah should just chase the larger antelope. If $s < \ell < 2s$, then there are two pure Nash equilibria, (L,S) and (S,L). These pay off quite well for both cheetahs — but how would two healthy cheetahs agree which should chase the smaller antelope? Therefore it makes sense to look for symmetric mixed equilibria.

If the first cheetah chases the large antelope with probability x , then the expected payoff to the second cheetah by chasing the larger antelope is

equal to

$$\frac{\ell}{2}x + (1-x)\ell,$$

and that arising from chasing the smaller antelope is

$$xs + (1-x)\frac{s}{2}.$$

A mixed Nash equilibrium arises at that value of x for which these two quantities are equal, because, at any other value of x , player *II* would have cause to deviate from the mixed strategy $(x, 1-x)$ to the better of the pure strategies available. We find that

$$x = \frac{2\ell - s}{\ell + s}.$$

This actually yields a symmetric mixed equilibrium.

Symmetric mixed Nash equilibria are of particular interest. It has been experimentally verified that in some biological situations, systems approach such equilibria, presumably by mechanisms of natural selection. We explain briefly how this might work. First of all, it is natural to consider symmetric strategy pairs, because if the two players are drawn at random from the same large population, then the probabilities with which they follow a particular strategy are the same. Then, among symmetric strategy pairs, Nash equilibria play a special role. Consider the above mixed symmetric Nash equilibrium, in which $x_0 = (2\ell - s)/(\ell + s)$ is the probability of chasing the large antelope. Suppose that a population of cheetahs exhibits an overall probability $x > x_0$ for this behavior (having too many greedy cheetahs, or every single cheetah being slightly too greedy). Now, if a particular cheetah is presented with a competitor chosen randomly from this population, then chasing the small antelope has a higher expected payoff to this particular cheetah than chasing the large one. That is, the more modest a cheetah is, the larger advantage it has over the average cheetah. Similarly, if the cheetah population is too modest on the average, i.e., $x < x_0$, then the more ambitious cheetahs have an advantage over the average. Altogether, the population seems to be forced by evolution to chase antelopes according to the symmetric mixed Nash equilibrium.

Example: the game of chicken. Two drivers speed head-on toward each other and a collision is bound to occur unless one of them chickens out at the last minute. If both chicken out, everything is okay (they both win 1). If one chickens out and the other does not, then it is a great success for the player with iron nerves (payoff= 2) and a great disgrace for the chicken (payoff= -1). If both players have iron nerves, disaster strikes (both lose some big value M).

We solve the game of chicken. Write C for the strategy of chickening out, D for driving forward. The pure equilibria are (C, D) and (D, C) . To

determine the mixed equilibria, suppose that player I plays C with probability x and D with probability $1 - x$. This presents player II with expected payoffs of $2x - 1$ if she plays C , and $(M + 2)x - M$ if she plays D . We seek an equilibrium where player II has positive weight on each of C and D , and thus one for which

$$2x - 1 = (M + 2)x - M.$$

That is, $x = 1 - 1/M$. The payoff for player to is $2x - 1$, which equals $1 - 2/M$. Note that, as M increases to infinity, this symmetric mixed equilibrium gets concentrated on (C, C) , and the expected payoff increases up to 1.

There is an abstract paradox here. We have a symmetric game with payoff matrices A and B that has a unique symmetric equilibrium with payoff γ . By replacing A and B by smaller matrices $\tilde{A}$ and $\tilde{B}$, we obtain a payoff $\tilde{\gamma}$ in a unique symmetric equilibrium that exceeds γ . This is impossible in zero-sum games.

However, if the decision of each player gets switched randomly with some small but fixed probability, then letting $M \rightarrow \infty$ does not yield total concentration on the strategy pair (C, C) .

Furthermore, this is again a game in which the possibility of a binding commitment increases the payoff. If one player rips out the stirring wheel and throws it out of the car, then he makes it impossible to chicken out. If the other player sees this and believes her eyes, then she has no other choice but to chicken out.

One more example. In the battle of sexes and the game of chicken, making a binding commitment pushed the game into a pure Nash equilibrium, and the nature of that equilibrium strongly depended on who managed to commit first. In the game of chicken, the payoff for the one who did not make the commitment was lower than the payoff in the unique mixed Nash equilibrium, while, in the battle of sexes, it was higher. Here is an example where there is no pure equilibrium, only a unique mixed one, and both commitment strategy pairs have the property that the player who did not make the commitment still gets the Nash equilibrium payoff.

II	C	D
I		
A	(6, -10)	(0, 10)
B	(4, 1)	(1, 0)

In this game, there is no pure Nash equilibrium (one of the players always prefers another strategy, in a cyclic fashion). For mixed strategies, if player I plays (A, B) with probabilities $(p, 1 - p)$, and player II plays (C, D) with probabilities $(q, 1 - q)$, then the expected payoffs are $f(p, q) = 1 + 3q - p + 3pq$ for I and $g(p, q) = 10p + q - 21pq$ for II . We easily get that the unique mixed equilibrium is $p = 1/21$ and $q = 1/3$, with payoffs 2 for I and $10/21$ for II . If I can make a commitment, then by choosing $p = 1/21 - \epsilon$ for

some small $\epsilon > 0$ he will make *II* choose $q = 1$, and the payoffs will be $4 + 2/21 - 2\epsilon$ for *I* and $10/21 + 11\epsilon$ for *II*. If *II* can make a commitment, then by choosing $q = 1/3 + \epsilon$ she will make *I* choose $p = 1$, and the payoffs will be $2 + 6\epsilon$ for *I* and $10/3 - 11\epsilon$ for *II*.

4.3 General sum games with $k \geq 2$ players

It does not make sense to talk about zero-sum games when there are more than two players. The notion of a Nash equilibrium, however, can be used in this context. We now describe formally the set-up of a game with $k \geq 2$ players. Each player i has a set S_i of pure strategies. If, for each $i \in \{1, \dots, k\}$, player i uses strategy $l_i \in S_i$, then player j has a payoff of $F_j(l_1, \dots, l_k)$, where we are given functions $F_j : S_1 \times S_2 \times \dots \times S_k \rightarrow \mathbb{R}$, for $j \in \{1, \dots, k\}$.

Example: an ecology game. Three firms will either pollute a lake in the following year, or purify it. They pay 1 unit to purify, but it is free to pollute. If two or more pollute, then the water in the lake is useless, and each firm must pay 3 units to obtain the water that they need from elsewhere. If at most one firm pollutes, then the water is usable, and the firms incur no further costs. Assuming that firm *III* purifies, the cost matrix is:

	II	Pu	Po
I			
Pu		(1,1,1)	(1,0,1)
Po		(0,1,1)	(3,3,3+1)

If firm *III* pollutes, then it is:

	II	Pu	Po
I			
Pu		(1,1,0)	(3+1,3,3)
Po		(3,3+1,3)	(3,3,3)

To discuss the game, we firstly introduce the notion of Nash equilibrium in the context of games with several players:

Definition 9 A pure Nash equilibrium in a k -person game is a set of pure strategies for each of the players,

$$(l_1^*, \dots, l_k^*) \in S_1 \times \dots \times S_k$$

such that, for each $j \in \{1, \dots, k\}$ and l_j ,

$$F_j(l_1^*, \dots, l_{j-1}^*, l_j, l_{j+1}^*, \dots, l_k^*) \leq F_j(l_1^*, \dots, l_{j-1}^*, l_j^*, l_{j+1}^*, \dots, l_k^*).$$

More generally, a **mixed Nash equilibrium** is a collection of k probability vectors $\tilde{X}^i$, each of length $|S_i|$, such that

$$F_j(\tilde{X}^1, \dots, \tilde{X}^{j-1}, X, \tilde{X}^{j+1}, \dots, \tilde{X}^k) \leq F_j(\tilde{X}^1, \dots, \tilde{X}^{j-1}, \tilde{X}^j, \tilde{X}^{j+1}, \dots, \tilde{X}^k),$$

for each probability vector X of length $|S_j|$. We have written:

$$F_j(X^1, X^2, \dots, X^k) := \sum_{l_1 \in S_1, \dots, l_k \in S_k} X^1(l_1) \dots X^k(l_k) F_j(l_1, \dots, l_k).$$

Definition 10 A game is **symmetric** if, for every $i_0, j_0 \in \{1, \dots, k\}$, there is a permutation π of the set $\{1, \dots, d\}$ such that $\pi(i_0) = j_0$ and

$$F_{\pi(i)}(l_{\pi(1)}, \dots, l_{\pi(k)}) = F_i(l_1, \dots, l_k).$$

For this definition to make sense, we are in fact requiring that the strategy sets of the players coincide.

We will prove the following result:

Theorem 10 (Nash's theorem) Every game has a Nash equilibrium.

Note that the equilibrium may be mixed.

Corollary 1 In a symmetric game, there is a symmetric Nash equilibrium.

Returning to the ecology game, note that the pure equilibria consist of each firm polluting, or one of the three firms polluting, and the remaining two purifying. We now seek mixed equilibria. Let x, y, z be the probability that firm *I, II, III* purifies, respectively. If firm *III* purifies, then its expected cost is $xy + x(1-y) + y(1-x) + 4(1-x)(1-y)$. If it pollutes, then the cost is $3x(1-y) + 3y(1-x) + 3(1-x)(1-y)$. If we want an equilibrium with $0 < z < 1$, then these two expected values must coincide, which gives the equation $1 = 3(x + y - 2xy)$. Similarly, assuming $0 < y < 1$ we get $1 = 3(x + z - 2xz)$, and assuming $0 < x < 1$ we get $1 = 3(y + z - 2yz)$. Subtracting the second equation from the first one we get $0 = 3(y - z)(1 - 2x)$. If $y = z$, then the third equation becomes quadratic in y , with two solutions, $y = z = (3 \pm \sqrt{3})/6$, both in $(0, 1)$. Substituting these solutions into the first equation, both yield $x = y = z$, so there are two symmetric mixed equilibria. If, instead of $y = z$, we let $x = 1/2$, then the first equation becomes $1 = 3/2$, which is nonsense. This means that there is no asymmetric equilibrium with at least two mixed strategies. It is easy to check that there is no equilibrium with two pure and one mixed strategy. Thus we have found all Nash equilibria: one symmetric and three asymmetric pure equilibria, and two symmetric mixed ones.

4.4 The proof of Nash's theorem

Recall Nash's theorem:

Theorem 11 For any general sum game with $k \geq 2$ players, there exists at least one Nash equilibrium.

To prove this theorem, we will use:

Theorem 12 (Brouwer's fixed point theorem) *If $K \subseteq \mathbb{R}^d$ is closed, convex and bounded, and $T : K \rightarrow K$ is continuous, then there exists $x \in K$ such that $T(x) = x$.*

Remark. We will prove this fixed point theorem later, but observe now that it is easy in dimension $d = 1$, when K is just a closed interval $[a, b]$. Defining $f(x) = T(x) - x$, note that $T(a) \geq a$ implies that $f(a) \geq 0$, while $T(b) \leq b$ implies that $f(b) \leq 0$. The intermediate value theorem assures the existence of $x \in [a, b]$ for which $f(x) = 0$, and thus, so that $T(x) = x$. Note also that each of the hypotheses on K in the theorem are required. Consider $T : \mathbb{R} \rightarrow \mathbb{R}$ given by $T(x) = x + 1$, as well as $T : (0, 1) \rightarrow (0, 1)$ given by $T(x) = x/2$, and also, $T : \{z \in \mathbb{C} : |z| \in [1, 2]\} \rightarrow \{z \in \mathbb{C} : |z| \in [1, 2]\}$ given by $T(x) = x \exp(i\pi/2)$.

Proof of Nash's theorem using Brouwer's theorem. Suppose that there are two players and the game is specified by payoff matrices $A_{m \times n}$ and $B_{m \times n}$ for players I and II . We will define a map $F : K \rightarrow K$ (with $K = \Delta_m \times \Delta_n$) from a pair of strategies for the two players to another such pair. Note firstly that K is convex, closed and bounded. Define, for $x \in \Delta_m$ and $y \in \Delta_n$,

$$c_i(x, y) = c_i = \max \left\{ \sum_j a_{ij} y_j - x^T A y, 0 \right\}.$$

Or, $c_i = \max \{A_i y - x^T A y, 0\}$, where A_i denotes the i th row of the matrix A . That is, c_i is equal to the gain for player I obtained by switching from strategy x to pure strategy i , if this gain is positive: otherwise, it is zero. Similarly, we define

$$d_j(x, y) = d_j = \max \{x^T B^{(j)} - x^T B y, 0\},$$

where $B^{(j)}$ denotes the j -th column of B . The quantities d_j have the same interpretation for player II as the c_i do for player I . We now define the map F ; it is given by $F(x, y) = (\tilde{x}, \tilde{y})$, where

$$\tilde{x}_i = \frac{x_i + c_i}{1 + \sum_{k=1}^m c_k}$$

for $i \in \{1, \dots, m\}$, and

$$\tilde{y}_j = \frac{y_j + d_j}{1 + \sum_{k=1}^n d_k}$$

for $j \in \{1, \dots, n\}$. Note that F is continuous, because c_i and d_j are. Applying Brouwer's theorem, we find that there exists $(x, y) \in K$ for which $(x, y) = (\tilde{x}, \tilde{y})$. We now claim that, for this choice of x and y , each $c_i = 0$

for $i \in \{1, \dots, m\}$, and $d_j = 0$ for $j \in \{1, \dots, n\}$. To see this, suppose, for example, that $c_1 > 0$. Note that the current payoff of player I is a weighted average $\sum_{i=1}^m x_i A_i y$. There must exist $i \in \{1, \dots, m\}$ for which $x^T A y \geq A_i y$, and $x_i > 0$. For this i , we have that $c_i = 0$. This implies that

$$\tilde{x}_i = \frac{x_i + c_i}{1 + \sum_{k=1}^m c_k} < x_i,$$

because $c_1 > 0$. That is, the assumption that $c_1 > 0$ has brought a contradiction.

We may repeat this argument for each $i \in \{1, \dots, m\}$, thereby proving that each $c_i = 0$. Similarly, each $d_j = 0$. We deduce that $x^T A y \geq A_i y$ for all $i \in \{1, \dots, m\}$. This implies that

$$x^T A y \geq x'^T A y$$

for all $x' \in \Delta_m$. Similarly,

$$x^T B y \geq x^T B y'$$

for all $y' \in \Delta_n$. Thus, (x, y) is a Nash equilibrium. $\square$

For $k > 2$ players, we still can consider the functions

$$c_i^{(j)}(x^{(1)}, \dots, x^{(k)}) \quad \text{for } i, j = 1, \dots, k,$$

where $x^{(j)} \in \Delta_{n(j)}$ is a mixed strategy for player j , and $c_i^{(j)}$ is the gain for player j obtained by switching from strategy $x^{(j)}$ to pure strategy i , if this gain is positive. The simple matrix notation for $c_i^{(j)}$ is lost, but the proof carries over.

We also stated that in a symmetric game, there is always a **symmetric Nash equilibrium**. This also follows from the above proof, by noting that the map F , defined from the k -fold product $\Delta_n \times \dots \times \Delta_n$ to itself, can be restricted to the diagonal

$$D = \{(x, \dots, x) \in \Delta_n^k : x \in \Delta_n\}.$$

The image of D under F is again in D , because, in a symmetric game, $c_i^{(1)}(x, \dots, x) = \dots = c_i^{(k)}(x, \dots, x)$ for all $i = 1, \dots, k$ and $x \in \Delta_n$. Then, Brouwer's fixed point theorem gives us a fixed point within D , which is a symmetric Nash equilibrium.

4.4.1 Some more fixed point theorems

We will discuss some fixed point theorems, beginning with:

Theorem 13 (Banach's fixed point theorem) *Let K be a complete metric space. Suppose that $T : K \rightarrow K$ satisfies $d(Tx, Ty) \leq \lambda d(x, y)$ for all $x, y \in K$, with $0 < \lambda < 1$ fixed. Then T has a unique fixed point in K .*

Note: Recall that a metric space is **complete** if each Cauchy sequence therein converges to a point in the space. Consider the metric space being a subset of $\mathbb{R}^d$ and the metric d being Euclidean distance if you are not familiar with these definitions.

Proof. Uniqueness of the fixed point: if $Tx = x$ and $Ty = y$, then

$$d(x, y) = d(Tx, Ty) \leq \lambda d(x, y).$$

Thus, $d(x, y) = 0$, and $x = y$. As for existence, given any $x \in K$, we define $x_n = Tx_{n-1}$ for each $n \geq 1$, setting $x_0 = x$. Set $a = d(x_0, x_1)$, and note that $d(x_n, x_{n+1}) \leq \lambda^n a$. If $k > n$, then

$$d(x_k, x_n) \leq d(x_n, x_{n+1}) + \dots + d(x_{k-1}, x_k) \leq a(\lambda^n + \dots + \lambda^{k-1}) \leq \frac{a\lambda^n}{1 - \lambda}.$$

This implies that $\{x_n : n \in \mathbb{N}\}$ is a Cauchy sequence. The metric space K is complete, whence $x_n \rightarrow z$ as $n \rightarrow \infty$. Note that

$$d(z, Tz) \leq d(z, x_n) + d(x_n, x_{n+1}) + d(x_{n+1}, Tz) \leq (1 + \lambda)d(z, x_n) + \lambda^n a \rightarrow 0$$

as $n \rightarrow \infty$. Hence, $d(Tz, z) = 0$, and $Tz = z$. $\square$

Example regarding Banach's fixed point theorem. Consider the map $T : \mathbb{R} \rightarrow \mathbb{R}$ given by

$$T(x) = x + \frac{1}{1 + \exp(x)}.$$

Note that, if $x < y$, then

$$T(x) - x = \frac{1}{1 + \exp(x)} > \frac{1}{1 + \exp(y)} = T(y) - y,$$

implying that $T(y) - T(x) < y - x$. Note also that

$$T'(x) = 1 - \frac{\exp(-x)}{(1 + \exp(x))^2} > 0,$$

so that $T(y) - T(x) > 0$. Thus, T decreases distances, but it has no fixed point. This is not a counterexample to Banach's fixed point theorem, however, because there does not exist any $\lambda \in (0, 1)$ for which $|T(x) - T(y)| < \lambda|x - y|$ for all $x, y \in \mathbb{R}$.

Theorem 14 (Compact fixed point theorem) *If X is a compact metric space and $T : X \rightarrow X$ satisfies $|T(x) - T(y)| < \lambda|x - y|$ for all $x, y \in X$, $x \neq y$, then T has a fixed point.*

Proof. Let $f : X \rightarrow X$ be given by $f(x) = d(Tx, x)$. Note that f is continuous:

$$|f(x) - f(y)| \leq |d(x, Tx) - d(y, Ty)| \leq d(x, y) + d(Tx, Ty) \leq 2d(x, y),$$

where the first inequality followed from two applications of the triangle inequality. By the compactness of X , there exists $x_0 \in X$ such that

$$f(x_0) = \inf_{x \in X} f(x). \quad (12)$$

If $Tx_0 \neq x_0$, then $d(x, Tx_0) > d(Tx_0, T^2x_0)$, contradiction (12). Thus implies that $Tx_0 = x_0$. $\square$

4.4.2 Sperner's lemma

We now state and prove a tool to be used in the proof of Brouwer's fixed point theorem.

Lemma 6 (Sperner) *In $d = 1$: suppose that the unit interval is subdivided $0 = t_0 < t_1 < \dots < t_n = 1$, with each t_i being marked zero or one, t_0 marked zero and t_1 marked one. Then the number of adjacent pairs (t_j, t_{j+1}) with different markings is odd. In $d = 2$: subdivide a triangle into smaller*




Figure 1: Sperner's lemma when $d = 2$.

triangles in such a way that a vertex of any of the small triangles may not lie in the interior of an edge of another. Label the vertices of the small triangles 0, 1 or 2: the three vertices of the big triangle must be labelled 0, 1 and 2; vertices of the small triangles that lie on an edge of the big triangle must receive the label of one of the endpoints of that edge. Then the number of small triangles with three differently labeled vertices is odd; in particular, it is non-zero.

Remark. Sperner's lemma holds in any dimension. In the general case d , we replace the triangle by a d -simplex, use d labels, with analogous restrictions on the labels used.

Proof. For $d = 1$, this is obvious. For $d = 2$, we will count in two ways the set Q of pairs of a small triangle and an edge on that triangle. Let A_{12} denote the number of 12 type edges in the boundary of the big triangle. Let B_{12} be the number of such edges in the interior. Let N_{abc} denote the number of triangles where the three labels are a , b and c . Note that

$$N_{012} + 2N_{112} + 2N_{122} = A_{12} + 2B_{12},$$

because each side of this equation is equal to the number of pairs of triangle and edge, where the edge is of type (12). From the case $d = 1$ of the lemma, we know that A_{12} is odd, and hence N_{012} is odd, too. (In general, we may induct on the dimension, and use the inductive hypothesis to find that this quantity is odd.) $\square$

Corollary 2 (No retraction theorem) *Let $K \subseteq \mathbb{R}^d$ be compact and convex, and with non-empty interior. There is no continuous map $f : K \rightarrow \partial K$ whose restriction to ∂K is the identity.*

Proof for $d = 2$. Firstly, we show that it suffices to take $K = \Delta$, where Δ is an equilateral triangle. For, we may locate $x \in K$ such that there exists a small triangle centered at x and contained in K , because K has a non-empty interior. We call this triangle Δ for convenience. Construct a map $h : K \rightarrow \Delta$ as follows. For each $y \in \partial K$, define $h(y)$ to be equal to that element of $\partial \Delta$ that the line segment from x through y intersects. Setting $h(x) = x$, define $h(z)$ for other $z \in K$ by a linear interpolation of the values $h(x)$ and $h(q)$, where q is the element of ∂K lying on the line segment from x through z .

Note that, if $F : K \rightarrow \partial K$ is a retraction from K to ∂K , then $h \circ F \circ h^{-1} : \Delta \rightarrow \partial \Delta$ is a retraction of Δ . This is the reduction we claimed.

Now suppose that $F_\Delta : \Delta \rightarrow \partial \Delta$ is a retraction of the equilateral triangle with sidelength 1. Since $F = F_\Delta$ is continuous and Δ is compact, there exists $\delta > 0$ such that $x, y \in \Delta$ satisfying $|x - y| < \delta$ also satisfy $|f(x) - f(y)| < \frac{\sqrt{3}}{4}$.

Label the three vertices of Δ by 0, 1, 2. Triangulate Δ into triangles of sidelength less than δ . In this subdivision, label any vertex x according to the label of the vertex of Δ nearest to $F(x)$, with an arbitrary choice being made to break ties.

By Sperner's lemma, there exists a small triangle whose vertices are labelled 0, 1, 2. The condition that $|f(x) - f(y)| < \frac{\sqrt{3}}{4}$ implies that any pair of these vertices must be mapped under F to interior points of one of the side of Δ , with a different side of Δ for each pair. This is impossible, implying that no retraction of Δ exists. $\square$

Note. As we already remarked in the Introduction, Brouwer's fixed point theorem fails if the convexity assumption is completely omitted. This is also true for the above corollary. However, the main property of K that we used was not convexity; it is enough if there is a homeomorphism (a 1-1 continuous map with continuous inverse) between K and Δ .

4.4.3 Proof of Brouwer's fixed point theorem

Recall that we are given a continuous map $T : K \rightarrow K$, with K a closed, bounded and convex set. Suppose that T has no fixed point. Then we can define a continuous map $F : K \rightarrow \partial K$ as follows. For each $x \in K$, we draw a ray from $T(x)$ through x until it meets ∂K . We set $F(x)$ equal to this point of intersection. If $T(x) \in \partial K$, we set $F(x)$ equal that intersection point of the ray with ∂K which is not equal to $T(x)$. E.g., in the case of the domain $K = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$, the map F may be written explicitly in terms of T . With some checking, it follows that $F : K \rightarrow \partial K$ is continuous. Thus, F is a retraction of K — but this contradicts the No Retraction Theorem, so T must have a fixed point. $\square$

4.5 Some further examples

Example: the fish-selling game. Fish being sold at the market is fresh with probability $2/3$ and old otherwise, and the customer knows this. The seller knows whether the particular fish on sale now is fresh or old. The customer asks the fish-seller whether the fish is fresh, the seller answers, and then the customer decides to buy the fish, or to leave without buying it. The price asked for the fish is \$12. It is worth \$15 to the costumer if fresh, and nothing if it is old. The seller bought the fish for \$6, and if it remains unsold, then he can sell it to another seller for the same \$6 if it is fresh, and he has to throw it out if it is old. On the other hand, if the fish is old, the seller claims it to be fresh, and the customer buys it, then the seller loses \$R in reputation.




Figure 2: The Kuhn tree for the fish-selling game.

The tree of all possible scenarios, with the net payoffs shown as (seller,customer), is depicted in the figure. This is called the Kuhn tree of the game.

The seller clearly should not say "old" if the fish is fresh, hence we should examine two possible pure strategies for him: "FF" means he always says "fresh"; "FO" means he always tells the truth. For the customer, there are four ways to react to what he might hear. Hearing "old" means that the fish is indeed old, so it is clear that he should leave in this case. Thus two rational strategies remain: BL means he buys the fish if he hears "fresh" and leaves if he hears "old"; LL means he just always leaves. Here are the expected payoffs for the two players, with randomness coming from the actual condition of the fish.

C	BL	LL
S		
"FF"	(6-R/3,-2)	(-2,0)
"FO"	(2,2)	(-2,0)

We see that if losing reputation does not cost too much in dollars, i.e., if $R < 12$, then there is only one pure Nash equilibrium: "FF" against LL. However, if $R \geq 12$, then the ("FO",BL) pair also becomes a pure equilibrium, and the payoff for this pair is much higher than for the other equilibrium.


4.6 Potential games

We now discuss a collection of games called **potential games**, which are k -player general sum games that have a special feature. Let $F_i(s_1, s_2, \dots, s_k)$ denote the payoff to player i if the players adopt the pure strategies $s_1, s_2, \dots, s_k$, respectively. In a potential game, there is a function $\psi : S_1 \times \dots \times S_k \rightarrow \mathbb{R}$, defined on the product of the players' strategy spaces, such that

$$\begin{aligned} & F_i(s_1, \dots, s_{i-1}, \tilde{s}_i, s_{i+1}, \dots, s_k) - F_i(s_1, \dots, s_k) \\ &= \psi(s_1, \dots, s_{i-1}, \tilde{s}_i, s_{i+1}, \dots, s_k) - \psi(s_1, \dots, s_k), \end{aligned} \quad (13)$$

for each i . We assume that each S_i is finite. We call the function ψ the **potential function** associated with the game.

Example of a potential game: a simultaneous congestion game. In this sort of game, the cost of using each road depends on the number of users of the road. For the road AC , it is l_i if there are i users, for $i \in \{1, 2\}$, in the case of the game depicted in the figure. Note that the cost paid by a given driver depends only on the number of users, not on which user she is.



More generally, for k drivers, we may define $\mathbb{R}$ -valued map C on the product space of the road-index set and the set $\{1, \dots, k\}$, so that $C(j, u_j)$ is equal to the cost incurred by any driver using road j in the case that the total number of drivers using this road is equal to u_j . Note that the strategy vector $s = (s_1, s_2, \dots, s_k)$ determines the usage of each road. That is, it determines $u_i(s)$ for each $i \in \{1, \dots, k\}$, where

$$u_i(s) = \left| \left\{ j \in \{1, \dots, k\} : \text{player } j \text{ uses road } i \text{ under strategy } s_j \right\} \right|,$$

for $i \in \{1, \dots, R\}$ (with R being the number of roads.)

In the case of the game depicted in the figure, we suppose that two drivers, I and II , have to travel from A to D , or from B to C , respectively.

In general, we set

$$\psi(s_1, \dots, s_k) = - \sum_{r=1}^R \sum_{l=1}^{u_r(s)} C(r, l).$$

We claim that ψ is a potential function for such a game. We show why this is so in the specific example. Suppose that driver 1, using roads 1 and 2, makes a decision to use roads 3 and 4 instead. What will be the effect on her cost? The answer is a change of

$$\left(C(3, u_3(s) + 1) + c_3(4, u_4(s) + 1) \right) - \left(C(1, u_1(s)) + C(2, u_2(s)) \right).$$

How did the potential function change as a result of her decision? We find that, in fact,

$$\psi(s) - \psi(\tilde{s}) = C(3, u_3(s) + 1) + c_3(4, u_4(s) + 1) - C(1, u_1(s)) - C(2, u_2(s))$$

where $\tilde{s}$ denotes the new joint strategy (after her decision), and s denotes the previous one. Noting that payoff is the negation of cost, we find that the change in payoff is equal to the change in the value of ψ . To show that ψ is indeed a potential function, it would be necessary to reprise this argument in the case of a general change in strategy by one of the players.

Now, we have the following result:

Theorem 15 (Monderer and Shapley, Rosenthal) *Every potential game has a Nash equilibrium in pure strategies.*

Proof. By the finiteness of the set $S_1 \times \dots \times S_k$, there exists some s that maximizes $\psi(s)$. Note that the expression in (13) is at most zero, for any $i \in \{1, \dots, k\}$ and any choice of $\tilde{s}_i$. This implies that s is a Nash equilibrium. $\square$

It is interesting to note that the very natural idea of looking for a Nash equilibrium by minimizing $\sum_{r=1}^R u_r C(r, u_r)$ does not work.

5 Coalitions and Shapley value

We will discuss two further topics in this class. One is that of finding evolutionarily stable strategies: which Nash equilibria arise naturally? We will discuss only some examples, of which the celebrated game of ‘hawks and doves’ is one.

The topic we now turn to is that of games involving coalitions. Suppose we have a group of $k > 2$ players. Each seeks a part of a given prize, but may achieve that prize only by joining forces with some of the other players. The players have varying influence — but how much power does each have? This is a pretty general summary. We describe the theory in the context of an example.

5.1 The Shapley value and the glove market

We discuss an example, mentioned in the Introduction. A customer enters a shop seeking to buy a pair of gloves. In the store are the three players. Player 1 has a left glove and players 2 and 3 each have a right glove. The customer will make a payment of 100 dollars for a pair of gloves. In their negotiations prior to the purchase, how much can each player realistically demand of the payment made by the customer? To resolve this question, we introduce a **characteristic function** v , defined on subsets of the player set. By an abuse of notation, we will write v_{12} in place of $v_{\{1,2\}}$, and so on. The function v will take the values 0 or 1, and will take the value 1 precisely when the subset of players in question are able between them to affect their aim. In this case, this means that the subset includes one player with a left glove, and one with a right one — so that, between them, they may offer the customer a pair of gloves. Thus, the values are

$$v_{123} = v_{23} = v_{13} = 1,$$

and the value is 0 on every other subset of $\{1, 2, 3\}$. Note that v is a $\{0, 1\}$ -valued monotone function: if $S \subseteq T$, then $v_S \leq v_T$. Such a function is always **superadditive**: $v(S \cup T) \geq v(S) + v(T)$ if S and T are disjoint.

In general, a characteristic function is just a superadditive function with $v(\emptyset) = 0$. Shapley was searching for a value function ψ_i , $i \in \{1, \dots, k\}$, such that $\psi_i(v)$ would be the **arbitration value** (now called **Shapley value**) for player i in a game whose characteristic function is v . Shapley analysed this problem by introducing the following axioms:

Symmetry: if $v(S \cup \{i\}) = v(S \cup \{j\})$ for all S with $i, j \notin S$, then $\psi_i(v) = \psi_j(v)$.

No power / no value: if $v(S \cup \{i\}) = v(S)$ for all S , then $\psi_i(v) = 0$.

Additivity: $\psi_i(v + u) = \psi_i(v) + \psi_i(u)$.

Efficiency: $\sum_{i=1}^n \psi_i(v) = v(\{1, \dots, n\})$.

The second one is also called the “dummy” axiom. The third axiom is the most problematic: it assumes that for any of the players, there is no effect of earlier games on later ones.

Theorem 16 (Shapley) *There exists a unique solution for ψ .*

A simpler example first: For a fixed subset $S \subseteq \{1, \dots, n\}$, consider the **S -veto game**, in which the effective coalitions are those that contain each member of S . This game has characteristic function w_S , given by $w_S(T) = 1$ if and only if $S \subseteq T$. It is easy to find the unique function that is a Shapley value. Firstly, the “dummy” axiom gives that

$$\psi_i(w_S) = 0 \quad \text{if } i \notin S.$$

Then, for $i, j \in S$, the “symmetry” axiom gives $\psi_i(w_S) = \psi_j(w_S)$. This and the “efficiency” axiom imply

$$\psi_i(w_S) = \frac{1}{|S|} \quad \text{if } i \in S,$$

and we have determined the Shapley value. Moreover, we have that $\psi_i(cw_S) = c\psi_i(w_S)$ for any $c \in [0, \infty)$.

Now, note that the glove market game has the same payoffs as $w_{23} + w_{13}$, except for the case of the set $\{1, 2, 3\}$. In fact, we have that

$$w_{23} + w_{13} = v + w_{123}.$$

In particular, the “additivity” axiom gives

$$\psi_i(w_{23}) + \psi_i(w_{12}) = \psi_i(v) + \psi_i(w_{123}).$$

If $i = 1$, then $0 + 1/2 = \psi_1(v) + 1/3$, while, if $i = 3$, then $1/2 + 1/2 = \psi_3(v) + 1/3$. Hence $\psi_3(v) = 2/3$ and $\psi_1(v) = \psi_2(v) = 1/6$. This means that player *I* has two-thirds of the arbitration value, while player *II* and *III* have one-third between them.

Example: the four stockholders. Four people own stock in ACME. Player i holds i units of stock, for each $i \in \{1, 2, 3, 4\}$. Six shares are needed to pass a resolution at the board meeting. How much is the position of each player worth in the sense of Shapley value? Note that

$$1 = v_{1234} = v_{24} = v_{34},$$

while $v = 1$ on any 3-tuple, and $v = 0$ in each other case.

We will assume that the value v may be written in the form

$$v = \sum_{\emptyset \neq S} c_S w_S.$$

Later, we will see that there always exists such a way of writing v . For now, however, we assume this, and compute the coefficients c_S . Note first that

$$0 = v_1 = c_1$$

(we write c_1 for $c_{\{1\}}$, and so on). Similarly,

$$0 = c_2 = c_3 = c_4.$$

Also,

$$0 = v_{12} = c_1 + c_2 + c_{12},$$

implying that $c_{12} = 0$. Similarly,

$$c_{13} = c_{14} = 0.$$

Next,

$$1 = v_{24} = c_2 + c_4 + c_{24},$$

implying that $c_{24} = 0$. Similarly, $c_{34} = 1$. We have that

$$1 = v_{123} = c_{123},$$

while

$$1 = v_{124} = c_{24} + c_{124},$$

implying that $c_{124} = 0$. Similarly, $c_{134} = 0$, and

$$1 = v_{234} = c_{24} + c_{34} + c_{123} + c_{124} + c_{134} + c_{234} + c_{1234} = 1 + 1 + 1 + 0 + 0 - 1 + c_{1234},$$

implying that $c_{1234} = -1$. Thus,

$$v = w_{24} + w_{34} + w_{123} - w_{234} - w_{1234},$$

whence

$$\psi_1(v) = 1/3 - 1/4 = 1/12,$$

and

$$\psi_2(v) = 1/2 + 1/3 - 1/3 - 1/4 = 1/4,$$

while $\psi_3(v) = 1/4$, by symmetry with player 2. Finally, $\psi_4(v) = 5/12$.

5.2 Probabilistic interpretation of Shapley value

Suppose that the players arrive at the board meeting in a uniform random order. Then there exists a moment when, with the arrival of the next stockholder, the coalition already present in the board-room becomes effective. The Shapley value of a given player is the probability of that player being the one to make the existing coalition effective. We will now prove this assertion.

Recall that we are given $v(S)$ for all sets $S \subseteq [n] := \{1, \dots, n\}$, with $v(\emptyset) = 0$, and $v(S \cup T) \geq v(S) + v(T)$ if $S, T \subseteq [n]$ are disjoint.

Theorem 17 *Shapley's four axioms uniquely determine the functions ϕ_i . Moreover, we have the random arrival formula:*

$$\phi_i(v) = \frac{1}{n!} \sum_{k=1}^n \sum_{\pi \in S_n: \pi(k)=i} \left(v(\pi(1), \dots, \pi(k)) - v(\pi(1), \dots, \pi(k-1)) \right)$$

Remark. Note that this formula indeed specifies the probability just mentioned.

Proof. Recall the game for which $w_S(T) = 1$ if $S \subseteq T$, and $w_S(T) = 0$ in the other case. We showed that $\phi_i(w_S) = 1/|S|$ if $i \in S$, and $\phi_i(w_S) = 0$ otherwise. Our aim is, given v , to find coefficients $\{c_S\}_{S \subseteq [n], S \neq \emptyset}$ such that

$$v = \sum_{\emptyset \neq S \subseteq [n]} c_S w_S. \quad (14)$$

Firstly, we will assume (14), and determine the values of $\{c_S\}$. Applying (14) to the singleton $\{i\}$:

$$v(\{i\}) = \sum_{\emptyset \neq S \subseteq [n]} c_S w_S(\{i\}) = c_{\{i\}} w_{\{i\}}(i) = c_i, \quad (15)$$

where we may write c_i in place of $c_{\{i\}}$. More generally, suppose that we have determined c_S for all S with $|S| < l$. We want to determine $c_{\tilde{S}}$ for some $\tilde{S}$ with $|\tilde{S}| = l$. We have that

$$v(\tilde{S}) = \sum_{\emptyset \neq S \subseteq [n]} c_S w_S(\tilde{S}) = \sum_{S \subseteq \tilde{S}, |S| < l} c_S + c_{\tilde{S}}. \quad (16)$$

This determines $c_{\tilde{S}}$. Now let us verify that (14) does indeed hold. Define the coefficients $\{c_S\}$ via (15) and (16), inductively for sets $\tilde{S}$ of size $l > 1$; that is,

$$c_{\tilde{S}} = v(\tilde{S}) - \sum_{s \subseteq \tilde{S}: |s| < l} c_s.$$

However, once (15) and (16) are satisfied, (14) also holds (something that should be checked by induction). We now find that

$$\phi_i(v) = \phi_i \left(\sum_{\emptyset \neq S \subseteq [n]} c_S w_S \right) = \sum_{\emptyset \neq S \subseteq [n]} \phi_i(c_S w_S) = \sum_{S \subseteq [n], i \in S} \frac{c_S}{|S|}.$$

This completes the proof of the first statement made in the theorem.

As for the second statement: for each permutation π with $\pi(k) = i$, we define

$$\psi_i(v, \pi) = v(\pi(1), \dots, \pi(k)) - v(\pi(1), \dots, \pi(k-1)),$$

and

$$\Psi_i(v) = \frac{1}{n!} \sum_{\pi} \psi_i(v, \pi).$$

Our goal is to show that $\Psi_i(v)$ satisfies all four axioms.

For a given π , note that $\psi_i(v, \pi)$ satisfies the “dummy” and “efficiency” axioms. It also satisfies the “additivity” axiom, but not the “symmetry” axiom. We now show that averaging produces a new object that is already symmetric — that is, that $\{\Psi_i(v)\}$ satisfies this axiom. To this end, suppose that i and j are such that

$$v(S \cup \{i\}) = v(S \cup \{j\})$$

for all $S \subseteq [n]$ with $S \cap \{i, j\} = \emptyset$. For every permutation π , define π^* that switches the locations of i and j . That is, if $\pi(k) = i$ and $\pi(l) = j$, then $\pi^*(k) = j$ and $\pi^*(l) = i$, with $\pi^*(r) = \pi(r)$ with $r \neq k, l$. We claim that

$$\psi_i(v, \pi) = \psi_j(v, \pi^*).$$

Suppose that $\pi(k) = i$ and $\pi(l) = j$. Note that $\psi_i(v, \pi)$ contains the term

$$v(\pi(1), \dots, \pi(k)) - v(\pi(1), \dots, \pi(k-1)),$$

whereas $\psi_i(v, \pi^*)$ contains the corresponding term

$$v(\pi^*(1), \dots, \pi^*(k)) - v(\pi^*(1), \dots, \pi^*(k-1)).$$

We find that

$$\begin{aligned} \Psi_i(v) &= \frac{1}{n!} \sum_{\pi \in S_n} \psi_i(v, \pi) = \frac{1}{n!} \sum_{\pi \in S_n} \psi_j(v, \pi^*) \\ &= \frac{1}{n!} \sum_{\pi^* \in S_n} \psi_j(v, \pi^*) = \Psi_j(v), \end{aligned}$$

where in the second equality, we used the fact that the map $\pi \mapsto \pi^*$ is a one-to-one map from S_n to itself, for which $\pi^{**} = \pi$. Therefore, $\Psi_i(v)$ is indeed the unique Shapley value function. $\square$

5.3 Two more examples

A fish without intrinsic value. A seller has a fish having no intrinsic value to him, i.e., he values it at zero dollars. A buyer values the fish at 10 dollars. We find the Shapley value: suppose that the buyer pays x for the fish, with $0 < x \leq 10$. Writing S and B for the seller and buyer, we have that $v(S) = 0$, $v(B) = 0$, with $v(S, B) = (10 - x) + x$, so that $\phi_s(v) = \phi_B(v) = 5$.

A potential problem with using the Shapley value in this case is the possibility that the buyer under-reports his desire for the fish to the party that arbitrates the transaction.

Many right gloves. Find the Shapley values for the following variant of the glove game. There are $n = r + 2$ players. Players 1 and 2 have left gloves. The remaining players each have a right glove. Note that $V(S)$ is equal to the maximal number of proper and disjoint pairs of gloves. In other words, $V(S)$ is equal to the minimum of the number of left, and of right, gloves held by members of S . Note that $\phi_1(v) = \phi_2(v)$, and $\phi_r(v) = \phi_3(v)$, for each $r \geq 3$. Note also that

$$2\phi_1(v) + r\phi_3(v) = 2,$$

provided that $r \geq 2$. For which permutations does the third player add value to the coalition already formed? The answer is the following orders:

$$13, 23, \{1, 2\}3, \{1, 2, j\}3,$$

where j is any value in $\{4, \dots, n\}$, and where the curly brackets mean that each of the resulting orders is to be included. The number of permutations corresponding to these possibilities is: $r!$, $r!$, $2(r-1)!$, and $6(r-1) \cdot (r-2)!$. This gives that

$$\phi_3(v) = \frac{2r! + 8(r-1)!}{(r+2)!}.$$

That is,

$$\phi_3(v) = \frac{2r + 8}{(r+2)(r+1)r}.$$

6 Mechanism design

Example: Keeping the meteorologist honest. The employer of a weatherman is determined that he should provide a good prediction of the weather for the following day. The weatherman's instruments are good, and he can, with sufficient effort, tune them to obtain the correct value for the probability of rain on the next day. There are many days, and, on the i -th of them, this correct probability is called p_i . On the evening of the $i - 1$ -th day, the weatherman submits his estimate $\hat{p}_i$ for the probability of rain on the following day, the i -th one. Which scheme should we adopt to reward or penalize the weatherman for his predictions, so that he is motivated to correctly determine p_i (that is, to declare $\hat{p}_i = p_i$)? The employer does not know what p_i is, because he has no access to technical equipment, but he does know the $\hat{p}_i$ values that the weatherman provides, and he knows whether or not it is raining on each day.

One suggestion is to pay the weatherman, on the i -th day, the amount $\hat{p}_i$ (or some dollar multiple of that amount) if it rains, and $1 - \hat{p}_i$ if it shines. If $\hat{p}_i = p_i = 0.6$, then the payoff is

$$\begin{aligned} \hat{p}_i \mathbb{P}(\text{it rains}) + (1 - \hat{p}_i) \mathbb{P}(\text{it shines}) &= \hat{p}_i p_i + (1 - \hat{p}_i)(1 - p_i) \\ &= 0.6 \times 0.6 + 0.4 \times 0.4 = 0.52. \end{aligned}$$

But in this case, even if the weatherman does correctly compute that $p_i = 0.6$, he is tempted to report the $\hat{p}_i$ value of 1, because, by the same formula, in this case, his earnings are 0.6.

Another idea is to wait for a long time, one year, say, and reward the weatherman according to how accurate his predictions have been on the average. More concretely, suppose for the sake of simplicity that the weatherman is only able to report $\hat{p}_i$ values on a scale of 0.1: so that he has eleven choices, namely $\{k/10 : k \in \{0, \dots, 10\}\}$. When a year has gone by, the days of that year may be divided into eleven types, according to the $\hat{p}_i$ -value that the weatherman declared. We would reward the weatherman if, for each of the eleven values of k , the fraction of days it rained in the class of days on which $\hat{p}_i$ was declared to be the value $k/10$, is actually close to $k/10$.

A scheme like this is quite reasonable and would certainly ensure that the weatherman tuned his instruments at the beginning of the year. However, it is not perfect. For example, ordinary random fluctuation mean that the weatherman will probably be slightly wrong as the end of the year approaches. It might be that it rained on 95 percent of the days on which the weatherman declared $\hat{p} = 0.9$, while, on those which he declared at 0.6, the average in reality has been 55 percent. Suppose that on the evening of one of the later days he sees that his instruments accurately predict 0.9. He knows that it is very likely to rain on the next day. But he is tempted to declare the next day at 0.6 — that is, to set $\hat{p}_i = 0.6$ for the i in question,

because doing so will (with probability 0.9) boost his 0.6 category and allow his 0.9 category to catch up with the downpour.

In fact, it is possible to design a scheme whereby we decide day-by-day how to reward the weatherman only on the basis of his declaration from the previous evening, without encountering the kind of problem that the last scheme had. Suppose that we pay $f(\hat{p}_i)$ to the weatherman if it rains, and $f(1 - \hat{p}_i)$ if it shines on day i . If $p_i = p$ and $\hat{p}_i = x$, then the expected payment made on day i is equal to

$$pf(x) + (1 - p)f(1 - x) = g_p(x).$$

(We are defining $g_p(x)$ to be this left-hand-side, because we are interested in how the payout is expected to depend on the prediction x of the weatherman on a given day where the probability of rain is p). Our aim is to reward the weatherman if his $\hat{p}_i$ equals p_i , in other words, to ensure that the expected payout is maximized when $x = p$. This means that the function $g_p : [0, 1] \rightarrow \mathbb{R}$ should satisfy $g_p(p) > g_p(x)$ for all $x \in [0, 1] - \{p\}$. By checking the derivative of $g_p(x)$, we see that $f(x) = \log x$ or $f(x) = -(1 - x)^2$ are good choices.